Linear Systems Analysis

by

R. K. Prasanth

Contents

Chapter 1

Models of Dynamic Systems

This course is about the analysis of models. A model is a mathematical representation of system behavior. It can be obtained by applying the laws of physics (*physics-based modeling*) or from experimental data (*black-box modeling*) or by combining both methods (*grey-box modeling*). Analysis deals with checking if a system model has a certain property. For example, *is the system stable ?*, or *how much is the overshoot in response to a unit step input ?*. We will define important system properties and provide computational procedures for analysis.

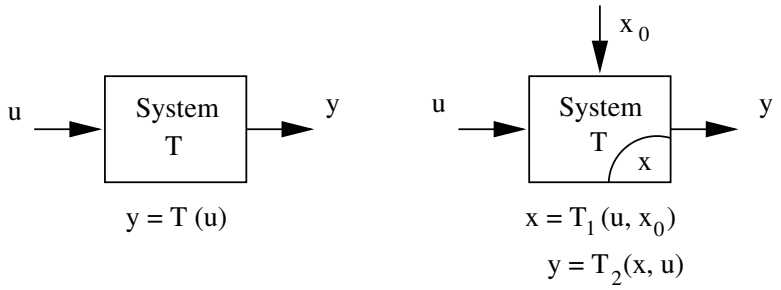

Figure 1.1: Input-output (left) and state space (right) representations

1.1 Input-output representations

One of the most powerful ways of representing a system is in terms of inputs and outputs. The simplest example of such a representation is the linear algebraic equation:

$$y = Tu$$

where T is a matrix and u is the *input* and y is the *output*. The equation specifies how the system (in this case the matrix T) acts on the input to produce the output. A more general case of algebraic input-output representation looks like:

$$Ay = Bu$$

where A and B are matrices. Extending further, when differential operators are involved, we get a dynamical system. For example,

$$\ddot{y} + a(t)\dot{y} + b(t)y = c(t)u \tag{1.1}$$

is a dynamical system represented in input-output form. If we denote differentiation with respect to time by $\mathcal{D}$, then the above equation can be written as:

$$\left(\mathcal{D}^2 + a(t)\mathcal{D} + b(t)\right) y = c(t)u$$

which has the same form as the linear algebraic equation. The coefficients a, b and c are referred to as system parameters.

An input-output representation relates signals that go into (*inputs* u) the box to the signals that come out (*outputs* y). The block diagram on the left hand side of Figure 1.1 shows this viewpoint. The mathematical formula contains only u, y and their derivatives. There is no explicit information about what is going on inside the box. This is the approach pioneered by engineers and scientists in the United States during the first half of 1900s.

1.2 State space representations

State space representations, on the other hand, gives us information about what goes on inside the box. They involve inputs, outputs and a set of variables called *internal states* (or simply states). Their general form is:

$$\dot{x} = f(x, u), \quad x(0) = x_0 \tag{1.2a}$$

$$y = h(x, u) \tag{1.2b}$$

where x is a vector of states, u is the input vector, y is the output vector and x_0 is the initial state. Inputs and outputs are physical quantities that are measured; so they do not depend on our particular point of view. On the other hand, the state of the box can be expressed using different sets of internal variables (or coordinates or states). For example, in a spring-mass-damper system, states could be the position x and the velocity $\dot{x}$ or they could be x and $x + 25\dot{x}$. Russian scientists and engineers pioneered this approach to systems.

The differential equations (1.2) in state space models are first order vector differential equations unlike the second order differential equation appearing in the input-output model (1.1). An ordinary differential equation of order n can always be converted into n first order differential equations by choosing states appropriately. For example, in (1.1), choose

$$x_1 = y \text{ and } x_2 = \dot{y}$$

as the states and form the state vector:

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Then,

$$\begin{aligned}\dot{x} &= \begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix}\\ &= \begin{bmatrix}\dot{y}\\ -a\dot{y} - by + cu\end{bmatrix} = \begin{bmatrix}x_2\\ -ax_2 - bx_1 + cu\end{bmatrix}\\ &= \begin{bmatrix}0 & 1\\ -b & -a\end{bmatrix}\begin{bmatrix}x_1\\ x_2\end{bmatrix} + \begin{bmatrix}0\\ c\end{bmatrix}u = Fx + Gu\end{aligned}$$

which gives us the differential equations in the state space representations. Now, for the output equation (1.2), we write:

$$y = [\,1 \quad 0\,]\,x$$

Therefore, only first order vector differential equations need to be considered.

We shall spend a great deal of time working with state space models, especially linear time-invariant (LTI) systems of the form:

$$\begin{aligned}\dot{x} &= Ax + Bu\\ y &= Cx + Du\end{aligned}$$

where A, B, C and D are matrices of appropriate dimensions. We shall often refer to the differential equation as the *state equation* and the algebraic equation for y as the *output equation*. There are three main motivations for studying this simple case: (i) all system properties in this simple case can be determined from linear algebraic properties of the matrices A, B, C and D that define the state space model, (ii) input-output and state space representations of a physical process are essentially equivalent which is very gratifying since physical properties should not depend on the modeler's choice of mathematics, and (iii) most real-world systems are nonlinear, but they behave like these linear systems in the vicinity of equilibrium points.

1.3 Some terminology

A system is *nonlinear* if at least one of the terms is a nonlinear function of the state variable. For example, the system:

$$\begin{aligned}\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} &= \begin{bmatrix}x_1^2 + tx_2\\ x_1\end{bmatrix}\\ y &= x_2\end{aligned}$$

is nonlinear even though the equations for $\dot{x}_2$ and y are linear. If at least one of the terms depends on time explicitly, then it is a *time-varying* equation. The above example is a time-varying nonlinear system.

If all the terms in a state space model are linear functions of the state, then the system is said to be *linear*. The following example is not a linear system:

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} = \begin{bmatrix}x_1 + 2x_2\\ x_1 - 2x_2\end{bmatrix}$$

$$y = \sin(x_1)$$

because the output equation is nonlinear even though the differential equation is linear. If the terms are not explicit functions of time, then the equation is *time-invariant*. The previous example is a nonlinear time-invariant system.

We will call the systems considered so far *homogeneous* as they do not contain any external inputs. External inputs can be control inputs (quantities that we have control over) or exogenous inputs (disturbances, command inputs etc). Systems with such inputs are *forced systems*. An example of a nonlinear time-varying forced system is the following:

$$\dot{x} = 5e^t/x + u$$

where x is the state and u is the external input.

1.4 Continuous-time and discrete-time

The independent variable time flows continuously in *continuous-time* systems. All systems mentioned in the previous sections and, in fact, models of most physical processes found in literature are continuous-time. *Discrete-time* systems are those in which time jumps discretely. Their general state space form is:

$$\begin{aligned} x(k+1) &= f(x(k), u(k)) \\ y(k) &= h(x(k), u(k)) \end{aligned}$$

where k denotes discrete-time and takes the values $0, 1, 2, \cdots$. The first equation above is called a *difference equation* and it specifies the states at the next time instant $k+1$. Note that the differential equation in the continuous time case specifies the rate of change of states. The general form of discrete-time linear time-invariant systems is:

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k \end{aligned}$$

where instead of the notation $x(k)$ we use x_k which is easier to write. It turns out that properties of discrete-time LTI systems can also be expressed in terms of the state space matrices (A, B, C, D).

Chapter 2

Introduction to Vector Spaces

Vector spaces give the algebraic structure needed to formulate and solve engineering problems. They form the foundation for linear algebra and optimization. This chapter gives a brief description of vector space theory.

2.1 Vector spaces and subspaces

A *field* $\mathbf{F}$ is a set together with the operations of addition and multiplication. The set must contain a unique additive inverse and a unique multiplicative inverse. The operations must be associative, commutative and distributive. Precise definition can be found in any good textbook on vector spaces or linear algebra [2, 3, 4]. Some examples of fields are the following.

Example 2.1.1 (Real and complex number fields) *The set of real numbers denoted by* $\mathbb{R}$ *along with the usual addition and multiplication is a field. The additive inverse is* 0 *and the multiplicative inverse is* 1. *The set of complex numbers* $\mathbb{C}$ *with the usual addition and multiplication is also a field.* $\triangle$

Example 2.1.2 (Rational number field) *The set of rational numbers along with the usual addition and multiplication is a field.* $\triangle$

The sets mentioned in the above examples have some nice features. They are *closed* under addition and multiplication. That is, in $\mathbb{R}$, addition (multiplication) of two real numbers always gives a real number. This algebraic feature[1] will be seen in later definitions. A second feature is that addition and multiplication are commutative, associative and distributive. In other words, the order in which addition and multiplication between elements are performed does not matter. This highly gratifying feature will also be seen later. A third feature is that, except for 0, all elements of these sets have a multiplicative inverse. For example, in

[1] Roughly speaking, like things produce like things under algebraic operations

$\mathbb{C}$, if $x \neq 0$ is a complex number, then there is a *complex number* y such that $xy = 1$. We normally write this inverse as $1/x$. Thus, the set of integers is not a field because the integer 2 has no inverse in the set of integers.

The elements of a field are called *scalars*. A field is the underlying object over which vector spaces are defined. The underlying field in most engineering problems is either $\mathbb{R}$ or $\mathbb{C}$(some engineering problems do have exotic fields).

Definition 2.1.1 (Vector space) *A vector (or linear) space over a field* $\mathbf{F}$ *is a triplet* $(V, +, \bullet)$*, where*

V *is a set of objects called vectors,*

$+$ *denotes vector space addition between elements of* V*, and*

$\bullet$ *denotes scalar multiplication of the elements of* V *by elements in the field* $\mathbf{F}$*,*

with the following properties:

1. V *is closed with respect to addition and scalar multiplication. That is, for any* x *and* y *in* V *and scalar* $\alpha \in \mathbf{F}$*, the vector sum* $x + y$ *and the product* $\alpha \bullet x$ *are in* V*.*

2. *Addition and scalar multiplication are associative, commutative and distributive. That is, for any* x, y, z *in* V *and* α, β *in* $\mathbf{F}$*, we have*

$$\begin{aligned}
(x + y) + z &= x + (y + z) && \textit{(addition is associative)} \\
x + y &= y + x && \textit{(addition is commutative)} \\
(\alpha\beta) \bullet x &= \alpha \bullet (\beta \bullet x) && \textit{(scalar multiplication is associative)} \\
\alpha \bullet (x + y) &= \alpha \bullet x + \alpha \bullet y && \textit{(addition distributes over scalar multiplication)} \\
(\alpha + \beta) \bullet x &= \alpha \bullet x + \beta \bullet x && \textit{(scalar multiplication distributes)}
\end{aligned}$$

3. *There exists a zero element, denoted by* 0*, in* V *such that*

$$x + 0 = x$$

for any x *in* V*.*

4. *For each* $x \in V$*, there exists an additive inverse, denoted by* $-x$ *such that* $x + (-x) = 0$*.*

5. *For each* $x \in V$*, we have* $1 \bullet x = x$*.* ★

The elements of V are called *vectors*, but they are not necessarily vectors in the colloquial sense. As we shall see in the examples below, they can be matrices, functions or other objects. The notation for scalar

multiplication $\bullet$ is tedious and is almost never used. For example, we normally write $2x$ instead of $2 \bullet x$ to mean the vector obtained by multiplying the vector x with the scalar 2.

Our motivation comes primarily from engineering applications where addition and scalar multiplication are defined in certain standard ways. For example, matrix addition is understood as addition of elements of the matrices, and function addition is performed point-wise. Such familiar notions will be called *standard or usual addition and scalar multiplication*. Whenever standard addition and scalar multiplication are used, we simply say that V is a vector space instead of the awkward notation $(V, +, \bullet)$. As mentioned, the most common fields used in engineering are $\mathbb{R}$ and $\mathbb{C}$. *A vector space is called a real vector space (respectively complex vector space) when the underlying field is* $\mathbb{R}$ *(respectively* $\mathbb{C}$*).*

Example 2.1.3 (Space of column vectors $\mathbf{F}^n$) *Consider the set* $\mathbf{F}^n$ *of column vectors of size* n *whose components are elements of the field* F*:*

$$\mathbf{F}^n \triangleq \left\{ x = \begin{bmatrix} x_1 \\ x_2 \\ \cdot \\ \cdot \\ x_n \end{bmatrix} : x_1, x_2, \cdots, x_n \text{ are in } \mathbf{F} \right\}$$

Given a column vector x*, the* i*th component of* x *is denoted by* x_i*. Let us define addition and scalar multiplication in the usual sense. That is, for vectors* $x, y \in \mathbf{F}^n$*, addition of* x *and* y *is defined componentwise:*

$$x + y = \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \cdot \\ \cdot \\ x_n + y_n \end{bmatrix}$$

and, scalar multiplication by $\alpha \in \mathbf{F}$ *is also defined componentwise:*

$$\alpha x = \begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \cdot \\ \cdot \\ \alpha x_n \end{bmatrix}$$

It is easy to verify that $\mathbf{F}^n$ *with this notion of addition and scalar multiplication is a vector space over* $\mathbf{F}$*. The key observation in the verification is that vector addition and scalar multiplication as defined above are* induced *from field addition and field multiplication (defined over* $\mathbf{F}$*). The properties of the latter induce the properties required of* $\mathbf{F}^n$ *to become a vector space.* $\triangle$

Example 2.1.4 (The spaces $\mathbb{R}$, $\mathbb{R}^n$, $\mathbb{C}$, $\mathbb{C}^n$) *Concrete examples of* $\mathbf{F}^n$ *can be found by taking the field* $\mathbf{F}$ *to be* $\mathbb{R}$ *or* $\mathbb{C}$*. Thus,* $\mathbb{R}$ *and* $\mathbb{R}^2$ *are examples of vector spaces over* $\mathbb{R}$*(or real vector spaces); while* $\mathbb{C}$ *and* $\mathbb{C}^2$ *are vector spaces over* $\mathbb{C}$*(or complex vector spaces).* $\triangle$

Example 2.1.5 (Spaces of matrices $\mathbf{F}^{m\times n}$, $\mathbb{R}^{m\times n}$, $\mathbb{C}^{m\times n}$) *Consider the set $\mathbf{F}^{m\times n}$ of $m \times n$ matrices whose components are elements of the field $\mathbf{F}$:*

$$\mathbf{F}^{m\times n} = \left\{ \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix} : a_{ij} \in \mathbf{F} \quad \forall i = 1, 2, \cdots, m,\ j = 1, 2, \cdots, n \right\}$$

Define addition and scalar multiplication in the usual component-wise sense: for matrices $x, y \in \mathbf{F}^{m\times n}$ and scalar $\alpha \in \mathbf{F}$,

$$x + y = [x_{ij} + y_{ij}]_{i,j=1}^{m,n}, \quad \alpha x = [\alpha x_{ij}]_{i,j=1}^{m,n}$$

where x_{ij} denotes the (i, j)th entry of x. Again, using the observation that these definitions are induced from field addition and field multiplication, we can show that $\mathbf{F}^{m\times n}$ is a vector space over $\mathbf{F}$.

As before, concrete examples are $\mathbb{R}^{m\times n}$ over $\mathbb{R}$ and $\mathbb{C}^{m\times n}$ over $\mathbb{C}$, respectively the real vector space of $m \times n$ matrices with real entries and the complex vector space of $m \times n$ matrices with complex entries. △

The above examples illustrate an important process of constructing new vector spaces by taking products of old ones. For example, the space $\mathbb{R}^3$ of real vectors of size 3 can be thought of as the Cartesian product of the spaces $\mathbb{R}^1$ and $\mathbb{R}^2$, or as the Cartesian product of $\mathbb{R}^1$ with itself three times. We state the process below and it is one of the ways of constructing complicated vector spaces from simpler ones. It enables us to "combine" and study vastly different objects and will be useful in advanced engineering problems.

Proposition 2.1.1 (Products of vector spaces) *Suppose that $(V, +_V, \bullet_V)$ and $(W, +_W, \bullet_W)$ are vector spaces over the field $\mathbf{F}$.The Cartesian product*

$$V \times W = \{(v, w) : v \in V \ \text{ and } \ w \in W\}$$

with the following notion of addition $+$ and scalar multiplication $\bullet$:

$$(v_1, w_1) + (v_2, w_2) = (v_1 +_V v_2,\ w_1 +_W w_2) \quad \text{for all } (v_1, w_1) \text{ and } (v_2, w_2) \text{ in } V \times W$$

$$\alpha \bullet (v, w) = (\alpha \bullet_V v,\ \alpha \bullet_W w) \quad \text{for all } \alpha \in \mathbf{F} \text{ and } (v, w) \text{ in } V \times W$$

is a vector space over $\mathbf{F}$. ■

Some remarks about the notation used in the proposition are in order. We represented the elements of the Cartesian product $V \times W$ as (v, w) where the *first component* v belongs to V and the *second component* w belongs to W. This is the standard representation. It has the advantage that the components can be different objects - vectors, matrices, functions, etc. Thus, we may consider the Cartesian product of $\mathbb{R}$ and $\mathbb{C}^{2\times 2}$:

$$\mathbb{R} \times \mathbb{C}^{2\times 2} = \left\{(x, y) : x \in \mathbb{R} \text{ and } y \in \mathbb{C}^{2\times 2}\right\}$$

whose typical element looks like:

$$\left(x, \begin{bmatrix} a & b \\ c & d \end{bmatrix}\right)$$

where x is a real number and a, b, c, d are complex numbers. When V and W are collections of objects of the same kind, it may be useful to represent elements of the Cartesian product as:

$$\begin{bmatrix} v \\ w \end{bmatrix}$$

For example, when $V = \mathbb{R}$ and $W = \mathbb{R}^2$, we may write:

$$\mathbb{R} \times \mathbb{R}^2 = \left\{ \begin{bmatrix} v \\ w_1 \\ w_2 \end{bmatrix} : v \in \mathbb{R} \text{ and } \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \in \mathbb{R}^2 \right\}$$

which suggests that the Cartesian product of the space of vectors of size 1 ($\mathbb{R}$) with the space of vectors of size 2 ($\mathbb{R}^2$) gives the space of vectors of size 3.

Another observation is regarding the notation used for addition and multiplication. V and W may have different notions of addition and scalar multiplication. We use subscripts to indicate the notional difference. Thus, $+_V$ and $\bullet_V$ denote addition and scalar multiplication in the vector space V. The definition of addition $+$ on the Cartesian product $V \times W$ is *induced* by the notions of additions in V and W. That is, the first components are added according to the addition $+_V$ in V and the second components are added according to $+_W$ in V. So, in general, the addition $+$ in $V \times W$ is an abstract composite notion. Similar comments hold for scalar multiplication.

Finally, the proposition deals with two vector spaces V and W. But, by repeatedly applying the proposition, we can get more complex spaces $U \times V \times W \times \cdots$. Indeed, the spaces $\mathbf{F}^n$ and $\mathbf{F}^{m \times n}$ defined earlier can be seen as applications of the above fact. For example, $\mathbf{F}^n$ is the product space of $\mathbf{F}$ taken n times.

Example 2.1.6 (Space of polynomials of degree at most n, $\mathcal{P}_n(T)$) *Let T be a subset of $\mathbb{R}$. A real polynomial p with domain T is a function of the form:*

$$p(t) = \sum_{k=0}^{m} p_k t^k$$

defined for all $t \in T$. p_k is called the kth polynomial coefficient. The coefficients of a polynomial are elements of some vector space. For example, a scalar polynomial is a polynomial whose coefficients are scalars, while a matrix-valued polynomial has coefficients in $\mathbf{F}^{m \times n}$. The zero polynomial *is the polynomial whose coefficients are all equal to zero. If at least one of the coefficients is non-zero, then the polynomial is refered to as a non-zero polynomial. We say that a non-zero polynomial p has degree n if t^n is the highest power with a non-zero coefficient (since the coefficients lie in a vector space, zero and non-zero are well-defined). The zero polynomial has degree zero.*

Let $n \geq 0$ be an integer. Let $\mathcal{P}_n(T)$ be the set of all polynomials with domain $T \subset \mathbb{R}$ of degree at most n whose coefficients are real numbers, i.e.

$$\mathcal{P}_n(T) \triangleq \left\{ \begin{array}{r} p : T \to \mathbb{R} \text{ such that } p(t) = \sum_{k=0}^{n} a_k t^k \text{ for all } t \in T \\ \text{and for some } a_0, a_1, \cdots, a_n \text{ in } \mathbb{R} \end{array} \right\} \tag{2.1}$$

Note that $\mathcal{P}_n(T)$ contains polynomials of degree $0, 1, \cdots, n$ as the coefficients are free variables and may take on the value zero.

Let us define addition and scalar multiplication of polynomials as follows:

$$\begin{aligned} (p+q)(t) &= (p(t) + q(t)) \qquad \forall\, t \in T \\ (\alpha p)(t) &= (\alpha p(t)) \qquad \forall\, t \in T \end{aligned}$$

for all p, q in $\mathcal{P}_n(T)$ and $\alpha \in \mathbb{R}$ (say in words the expressions on both sides of equality to fully understand the definitions).

With these definition, $\mathcal{P}_n(T)$ is a vector space over $\mathbb{R}$. △

$\mathcal{P}_n(T)$ is a fairly important vector space. It appears in many applications including curve fitting (i.e, find a polynomial degree at most n that fit a given data set). It is also our first example of a space of functions.

Example 2.1.7 (Space of polynomials $\mathcal{P}(T)$) *As before, let T be a subset of $\mathbb{R}$ and $\mathcal{P}_n(T)$ be the real vector space of polynomials defined on T of degree at most n whose coefficients are real numbers. Define*

$$\mathcal{P}(T) \triangleq \mathcal{P}_0(T) \bigcup \mathcal{P}_1(T) \bigcup \mathcal{P}_2(T) \cdots = \bigcup_{n=0}^{\infty} \mathcal{P}_n(T) \tag{2.2}$$

which is the set of all real polynomials. Define addition and scalar multiplication on $\mathcal{P}(T)$ in the standard way, i.e, point-wise as in Example 2.1.6. Then, $\mathcal{P}(T)$ is a real vector space. △

The difference between $\mathcal{P}_n(T)$ and $\mathcal{P}(T)$ lies in the bound on the polynomial degree. All the polynomials in $\mathcal{P}_n(T)$ are of degree at most n. That is, the positive integer n is a tight upper bound on the degree of the polynomials in $\mathcal{P}_n(T)$. On the other hand, there is no bound on the degree of polynomials in $\mathcal{P}(T)$. That is, given any positive integer m, there is a polynomial of degree strictly greater than m in $\mathcal{P}(T)$. Later on, we will introduce the concept of dimension of a vector space. It will become clear then that $\mathcal{P}_n(T)$ is a finite dimensional function space while $\mathcal{P}(T)$ is infinite dimensional.

Example 2.1.8 (Space of continuous functions $\mathcal{C}([a,b])$) *The set of real continuous functions defined on the interval (a, b) along with pointwise addition and scalar multiplication by real numbers is a real vector space. It is denoted by $\mathcal{C}([a,b])$.* △

Example 2.1.9 (Space of solutions of linear differential equations) *Consider the differential equation*

$$\dot{x} = Ax, \quad x(0) = x_0$$

where A is in $\mathbb{R}^{n\times n}$ and $x_0 \in \mathbb{R}^n$ is the initial condition. Let

$$\phi_{x_0} : [0, \infty) \to \mathbb{R}^n$$

denote the solution starting at x_0. Define:

$$V \stackrel{\Delta}{=} \{\phi_{x_0} : x_0 \in \mathbb{R}^n\}$$

as the set of all solutions. Then, V is a real vector space with the standard addition and scalar multiplication of functions. △

Let us consider the differential equation:

$$\dot{x} = -x, \quad x(0) = x_0$$

whose solution:

$$x(t) = e^{-t}x_0 \quad \text{for all } t \geq 0$$

is a continuous function (for a fixed initial condition x_0). So, the set of all solutions V consists of continuous functions on the time interval $[0, \infty)$. Therefore, V is a subset of $\mathcal{C}([0, \infty))$. Note that V and $\mathcal{C}([0, \infty))$ are both vector spaces. They also have the same notions of addition and scalar multiplication. This observation leads to the definition of a subspace which is as important as a vector space.

Definition 2.1.2 (Subspace) *A subset U of a vector space V is a subspace of V if and only if U is itself a vector space with the addition and scalar multiplication defined on V.* ⋆

Example 2.1.10 (Trivial subspaces) *The sets $\{0\}$ and V are the smallest and the largest (in a sense that shall become precise when we introduce dimension) subspaces of V.* △

Example 2.1.11 (Polynomials of degree at most n) *Let $\mathcal{P}(T)$ be the real vector space of polynomials given in Example 2.1.7. Let $n \geq 0$ be an integer. Consider the real vector space $\mathcal{P}_n(T)$ of polynomials of degree at most n defined in Example 2.1.6. Clearly, $\mathcal{P}_n(T)$ is a subset of $\mathcal{P}(T)$ as the latter contains polynomials of degree any degree. $\mathcal{P}_n(T)$ and $\mathcal{P}(T)$ have the same standard notions of addition and scalar multiplication, namely point-wise addition and scalar multiplication of functions. So, $\mathcal{P}_n(T)$ is a subspace of $\mathcal{P}(T)$.* △

To check if a subset U is a subspace, we may check if U is a vector space. This can be tedious as a number of properties need to be verified. The following result gives a simple test.

Proposition 2.1.2 (Subspace test) *Let U be a subset of a vector space V over the field $\mathbf{F}$. The following statements are equivalent.*

1. *U is a subspace of V*
2. *For any $x \in U$, $y \in U$, $\alpha \in \mathbf{F}$ and $\beta \in \mathbf{F}$, we have $\alpha x + \beta y \in U$.* ■

When $\alpha = 0$ and $\beta = 0$, the quantity

$$\alpha x + \beta y$$

evaluates to 0. So, *a subspace must contain* 0.

Example 2.1.12 (Continuous functions with a common zero) *Let $\mathcal{C}(\mathbb{R})$ be the real vector space of continuous functions on $\mathbb{R}$ with the standard function addition and scalar multiplication. Define:*

$$U \triangleq \{f \in \mathcal{C}(\mathbb{R}) : f(0) = 0\}$$

as the subset of all continuous functions that evaluate to zero at 0.

We claim that U is a subspace of $\mathcal{C}(\mathbb{R})$. To prove this, note that U is a subset of $\mathcal{C}(\mathbb{R})$ and that the function 0 is in U (0 is the function that is zero at every point in the domain). Now, pick any f and g in U and, α and β in $\mathbb{R}$. Define the function:

$$h = \alpha f + \beta g$$

Since f and g are in $\mathcal{C}(\mathbb{R})$ and $\mathcal{C}(\mathbb{R})$ is a vector space, h is in $\mathcal{C}(\mathbb{R})$. Moreover,

$$h(0) = \alpha f(0) + \beta g(0) = 0$$

as $f(0) = 0$ and $g(0) = 0$ based on the fact that f and g were chosen from U. So, h is in U. As these arguments apply to any f and g in U and, α and β in $\mathbb{R}$, we conclude using statement 2 of Proposition 2.1.2 that U is a subspace. △

Example 2.1.13 (Affine set) *Let U be a subspace of a vector space V and b be an element of V. Define*

$$A_b \triangleq \{y = x + b : x \in U\}$$

as the set of elements of V obtained by translating elements of U by a fixed amount b. The subset A_b is called an affine subset of V. When $b \in U$, A_b is a subspace of V. Otherwise, it is not. △

There are two ways of combining subspaces to produce new subspaces. They are discussed below. Recall that a subspace is a vector space by definition. So, the procedures to construct new subspaces given below can also be seen as ways of constructing new vector spaces from old ones.

Proposition 2.1.3 (Intersection of subspaces) *Let U and V be subspaces of a vector space W. Then, their intersection:*

$$U \bigcap V \triangleq \{x : x \in U \text{ and } x \in V\}$$

is a subspace of U, V and W. ■

The proof involves the application of Proposition 2.1.2. The expression on the right hand side of the equation in the above proposition involves no condition that is specific to vector spaces. In fact, it is the definition of intersection of two sets. Unlike intersection, the set-union of subspaces does not necessarily produce a subspace. We need a different notion of "union" of subspaces.

Definition 2.1.3 (Sum of subspaces) *Let U and V be subspaces of a vector space W. The sum of U and V is given by:*

$$U + V = \{z \in W : \text{ there exist } x \in U \text{ and } y \in V \text{ such that } z = x + y\}$$

If $U \bigcap V = \{0\}$ i.e., if zero is the only element in the intersection of the subspaces U and V, then the sum of U and V is called the direct sum of U and V. The direct sum is denoted by $U \oplus V$. ⋆

Proposition 2.1.4 (Sum and Direct Sum are subspaces) *Let U and V be subspaces of a vector space W. Then, their sum $U + V$ and their direct sum $U \oplus V$ are both subspaces of W.* ■

Example 2.1.14 *Let W be the vector space $\mathbb{R}^3$, the space of column vectors of size 3 with real entries. Let*

$$U \triangleq \left\{ \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : x_1 \in \mathbb{R},\ x_2 = 0,\ x_3 = 0 \right\} = \left\{ \begin{bmatrix} x_1 \\ 0 \\ 0 \end{bmatrix} \in \mathbb{R}^3 : x_1 \in \mathbb{R} \right\}$$

and

$$V \triangleq \left\{ \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 : x_1 = 0,\ x_2 \in \mathbb{R},\ x_3 = 0 \right\} = \left\{ \begin{bmatrix} 0 \\ x_2 \\ 0 \end{bmatrix} \in \mathbb{R}^3 : x_2 \in \mathbb{R} \right\}$$

It is easy to verify that U and V are subspaces of $W = \mathbb{R}^3$. Their direct sum is

$$\begin{aligned} U \oplus V &\triangleq \{z \in W : \text{ there exist unique elements } x \in U \text{ and } y \in V \text{ such that } z = x + y\} \\ &= \left\{ z \in W : \text{ there exist } x_1 \in \mathbb{R}, y_2 \in \mathbb{R} \text{ such that } z = \begin{bmatrix} x_1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ y_2 \\ 0 \end{bmatrix} \right\} \\ &= \left\{ \begin{bmatrix} x_1 \\ y_2 \\ 0 \end{bmatrix} : x_1 \in \mathbb{R}, y_2 \in \mathbb{R} \right\} \end{aligned}$$

It is not difficult to see that $U \oplus V$ is a subspace. △

An interesting property of direct sum of subspaces is given in the next proposition.

Proposition 2.1.5 (Property of Direct Sum) *Let U, V and W be subspaces of a vector space X. $W = U \oplus V$ if and only if*

$$W = \{z \in X : \text{ there exist unique elements } x \in U \text{ and } y \in V \text{ such that } z = x + y\}$$

■

That is, every vector in the direct sum of U and V can be written as the sum of a unique element in U and a unique element in V.

2.2 Linear combinations, independence, span and basis

This section introduces some important concepts that are used in many areas including the study of linear equations and optimization.

Definition 2.2.1 (Linear combination) *Let $\{x_k\}_{k=1}^m$ be a set of vectors in a vector space V (over $\mathbf{F}$). A linear combination of $\{x_k\}_{k=1}^m$ is an element x of V that can be expressed as:*

$$x = \sum_{k=1}^{m} \alpha_k x_k \tag{2.3}$$

for some scalars $\alpha_1, \alpha_2, \cdots, \alpha_m$ in $\mathbf{F}$. ★

A linear combination can be thought of as a weighted sum of vectors. Note that any vector of the weighted form (2.3) is in V because V is a vector space. So, as the weights $\{\alpha_k\}$ are varied, possibly different elements of V are generated. The set of all elements generated by this process is very special and is defined below.

Definition 2.2.2 (Span of $\{x_k\}_{k=1}^m$) *The span of a set of vectors $\{x_k\}_{k=1}^m$ in a vector space V is:*

$$\begin{aligned} \mathbf{Span}\left(\{x_k\}_{k=1}^m\right) &\triangleq \{x \in V : x \text{ is a linear combination of } \{x_k\}_{k=1}^m\} && \text{(2.4a)} \\ &= \left\{x \in V : x = \sum_{k=1}^{m} \alpha_k x_k \text{ for some } \alpha_1, \alpha_2, \cdots, \alpha_m \text{ in } \mathbf{F}\right\} && \text{(2.4b)} \end{aligned}$$

that is, the set of all vectors in V that are linear combinations of $\{x_k\}_{k=1}^m$. ★

It is instructive to write the weighted sum representation of a linear combination in the following way when V is $\mathbb{R}^n$ (or $\mathbb{C}^n$). Let $x_1, x_2, \cdots, x_m$ be vectors in $\mathbb{R}^n$. Define

$$X = [x_1 \;\; x_2 \;\; \cdots \;\; x_m]$$

as the $n \times m$ matrix obtained by placing the vectors side by side. So, the kth column of X is the vector x_k. Then, a linear combination x of $\{x_k\}_{k=1}^m$ can be written as:

$$x = \sum_{k=1}^{m} \alpha_k x_k = \sum_{k=1}^{m} x_k \alpha_k = X \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \cdot \\ \cdot \\ \alpha_m \end{bmatrix}}_{\alpha} = X\alpha$$

which shows that x is obtained by multiplying the matrix X by a vector α. From this view point, the span of $\{x_k\}_{k=1}^m$ is the set of all vectors that can be generated by multiplying X by vectors. We will use this observation later on to study matrices and systems of linear equations.

Definition 2.2.3 (Linear independence) *Let $\{x_k\}_{k=1}^m$ be a set of elements of a vector space V. The set $\{x_k\}_{k=1}^m$ is linearly independent if and only if*

$$\sum_{k=1}^{m} \alpha_k x_k = 0$$

for some scalars $\alpha_1, \alpha_2, \cdots, \alpha_m$ implies that $\alpha_1 = \alpha_2 = \cdots = \alpha_m = 0$. Otherwise, the set $\{x_k\}_{k=1}^m$ is linearly dependent. ★

Another way of stating linear independence is that the only set of scalars $\{\alpha_k\}$ that satisfy the equation:

$$\sum_{k=1}^{m} \alpha_k x_k = 0$$

is $\alpha_1 = \alpha_2 = \cdots = \alpha_m = 0$. When the set is linearly dependent, the equation has a solution with at least one non-zero α_k. Recall the definition of a linear combination of $\{x_k\}_{k=1}^m$ and the discussion following it using the matrix X whose columns are x_k's. Linear independence means that the linear system of equations

$$X\alpha = 0$$

has one and only one solution, namely $\alpha = 0$. Equivalently, linear dependence means that this system of equations has non-zero solutions.

Example 2.2.1 *Consider the vectors*

$$x_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \text{ and } x_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

in $\mathbb{R}^2$. The span of x_1 and x_2 is:

$$\mathbf{Span}\,(x_1, x_2) \;=\; \left\{ \alpha \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \beta \begin{bmatrix} 1 \\ -1 \end{bmatrix} : \alpha, \beta \in \mathbb{R} \right\}$$

$$
\begin{aligned}
&= \left\{ \begin{bmatrix} \alpha+\beta \\ \alpha-\beta \end{bmatrix} : \alpha, \beta \in \mathbb{R} \right\} \\
&= \left\{ \begin{bmatrix} \gamma \\ \delta \end{bmatrix} : \gamma, \delta \in \mathbb{R} \right\} \\
&= \mathbb{R}^2
\end{aligned}
$$

These computations can be explained very clearly using Figure 2.1 where x_1 and x_2 are shown by the directed arrows. The vectors x_1 and x_2 are perpendicular to each other. So, every vector in the plane can be expressed in terms of x_1 and x_2. Hence, the span is $\mathbb{R}^2$.

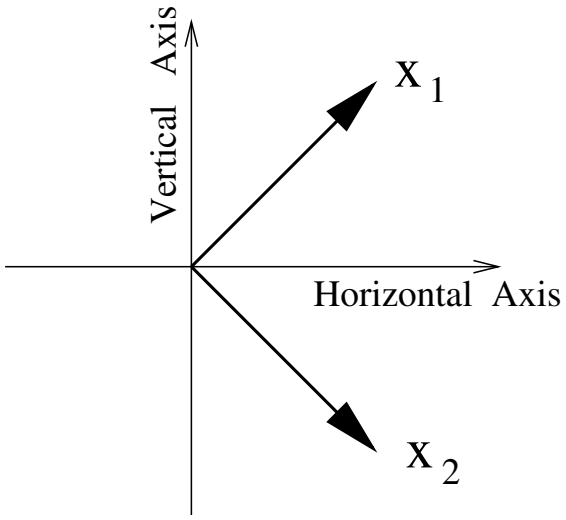

Figure 2.1: Representation of x_1 and x_2 of Example 2.2.1

Do the two vectors form a linearly independent set ? To answer this question, we appeal to the definition of linear independence. Accordingly, we must examine the condition:

$$\alpha_1 x_1 + \alpha_2 x_2 = 0$$

and must conclude that $\alpha_1 = 0$ and $\alpha_2 = 0$. Let us expand out the left hand side of the above equation:

$$\alpha_1 x_1 + \alpha_2 x_2 = \alpha_1 \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \alpha_2 \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \begin{bmatrix} \alpha_1 + \alpha_2 \\ \alpha_1 - \alpha_2 \end{bmatrix}$$

which shows that the condition to be examined is:

$$\begin{bmatrix} \alpha_1 + \alpha_2 \\ \alpha_1 - \alpha_2 \end{bmatrix} = 0$$

The above equation gives: $\alpha_1 + \alpha_2 = 0$ and $\alpha_1 - \alpha_2 = 0$ which holds if and only if $\alpha_1 = 0$ and $\alpha_2 = 0$. Therefore, we conclude that $\{x_1, x_2\}$ is a linearly independent set. △

Example 2.2.2 *The matrices*

$$x_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \; x_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \; x_3 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

form a linearly independent set in $\mathbb{R}^{2\times 2}$. *To see this, suppose that*

$$\alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 = 0$$

for some scalars $\alpha_1, \alpha_2, \alpha_3$. *Then, we have:*

$$\begin{bmatrix} \alpha_1 & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & \alpha_2 \end{bmatrix} + \begin{bmatrix} 0 & \alpha_3 \\ \alpha_3 & 0 \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_3 \\ \alpha_3 & \alpha_2 \end{bmatrix} = 0$$

which gives $\alpha_1 = 0$, $\alpha_2 = 0$ *and* $\alpha_3 = 0$ *as required. What is the span of these three matrices ?* △

Proposition 2.2.1 (Span is a subspace) *Let* $\{x_k\}_{k=1}^m$ *be a set of elements of a vector space* V. *Then, their span is a subspace of* V. ■

In view of this proposition, span is often referred to as the subspace spanned or generated by the set $\{x_k\}_{k=1}^m$. Span has a very nice property. Suppose a new vector x_{m+1} is added to the collection $\{x_k\}_{k=1}^m$. Then, the span of the new collection $\{x_k\}_{k=1}^{m+1}$ is at least as large as the span of $\{x_k\}_{k=1}^m$, if not bigger. This is stated precisely in the following result.

Proposition 2.2.2 (Nested sequence) *Let* $x_1, x_2, \cdots, x_{m+1}$ *be elements of a vector space* V. *The following statements are true.*

1. $\mathbf{Span}\left(\{x_k\}_{k=1}^m\right)$ *is contained in* $\mathbf{Span}\left(\{x_k\}_{k=1}^{m+1}\right)$
2. *Suppose that* $\{x_k\}_{k=1}^m$ *is linearly independent. Then,*

$$\mathbf{Span}\left(\{x_k\}_{k=1}^m\right) = \mathbf{Span}\left(\{x_k\}_{k=1}^{m+1}\right)$$

if and only if $\{x_k\}_{k=1}^{m+1}$ *is linearly dependent.* ■

Statement 1 says that

$$\mathbf{Span}\,(x_1) \subset \mathbf{Span}\,(x_1, x_2) \subset \cdots \mathbf{Span}\left(\{x_k\}_{k=1}^m\right) \subset \mathbf{Span}\left(\{x_k\}_{k=1}^{m+1}\right) \subset V$$

which shows the increasing (nested) property of span. Statement 2 gives a condition under which the inclusion becomes an equality.

The definitions of linear combination, span and independence involve a *finite* number of vectors $\{x_k\}_{k=1}^m$. Their extensions to any collection of vectors is needed in the sequel. We adopt the following extensions to avoid potential technical problems.

Definition 2.2.4 (Case of possibly infinite collection) *Let* U *be a subset of a vector space* V.

1. *A linear combination of elements of U is an element x in V which can be expressed as*
$$x = \sum_{k=1}^{m} \alpha_k x_k$$
for some finite collection $x_1, x_2, \cdots, x_m$ in U and scalars $\alpha_1, \alpha_2, \cdots, \alpha_m$.

2. *The elements of U are linearly independent (or, U is a linearly independent set) if, for any finite collection $\{x_k\}_{k=1}^m$ in U,*
$$\sum_{k=1}^{m} \alpha_k x_k = 0$$
implies $\alpha_1 = \alpha_2 = \cdots = \alpha_m = 0$.

3. *The span of elements of U is the set of all linear combinations of elements of U.* ⋆

These definitions clearly reduce to the earlier definitions when U is finite. The important point is that only a finite number of elements appear in the weighted sum even when U is an infinite set. Now, let us return to the properties of span. Proposition 2.2.1 says that the span is a subspace. This prompts us to ask if a subspace is the span of a set of vectors.

Proposition 2.2.3 (Subspace as span) *Let U be a subspace of a vector space V. Then, U is the span of elements of U.* ■

The proposition simply says that every element of a subspace can be expressed as a linear combination of elements of the subspace. This is not very useful in an operational sense because a subspace in general has an uncountable number of elements. In engineering applications, we will need to manipulate and represent subspaces. It would be nice if a subspace can be written as the span of a countable or, better yet, finite number of elements. These elements would then capture all the information contained in the subspace in a compressed form. In fact, the problems of data and image compression are precisely the problem of finding such elements. We introduce some relevant concepts and results.

Definition 2.2.5 (Dimension of spaces) *Let U be a subspace of a vector space V.*

1. *U is said to be finite dimensional if and only if U is the span of a finite number of elements in U. Otherwise, U is said to be infinite dimensional.*

2. *When U is finite dimensional, the minimum number of elements needed to span U is called the dimension of U.*

3. *When U is infinite dimensional, the dimension of U is ∞.* ■

Definition 2.2.6 (Basis) *Let U be a subspace of a vector space V. A basis for U is a linearly independent set of elements in U whose span is U.* ⋆

Theorem 2.2.1 (Basis and dimension) *Let U be a subspace of a vector space V. The following statements are true.*

1. *Let $\mathcal{B}$ be a basis for U. The dimension of U is equal to the number of elements in $\mathcal{B}$.*
2. *Suppose that U is finite dimensional. A subset $\mathcal{B}$ of U is a basis for U if and only if U is the span of $\mathcal{B}$ and the number of elements in $\mathcal{B}$ is equal to the dimension of U.* ■

The dimension of U is the minimum number of elements needed to span U. So, statement 1 of the theorem implies that a basis for U has the minimum number of elements among all subsets that span U. Put differently, if a collection has more elements than the dimension of its span, then the collection is not a basis. Statement 2 of the theorem is a special case of statement 1. A consequence is that every basis for a space has the same number of elements (more precisely, there is a one-to-one and onto mapping between different bases for a space).

Example 2.2.3 (Standard basis for $\mathbb{R}^n$, $\mathbb{C}^n$) *Let e_k be the $n \times 1$ vector whose entries are all zeros except for the kth element which is 1:*

$$e_1 = \begin{bmatrix} 1 \\ 0 \\ . \\ . \\ 0 \\ 0 \end{bmatrix}, \; e_2 = \begin{bmatrix} 0 \\ 1 \\ . \\ . \\ 0 \\ 0 \end{bmatrix}, \; \cdots, \; e_n = \begin{bmatrix} 0 \\ 0 \\ . \\ . \\ 0 \\ 1 \end{bmatrix}$$

The set of vectors $\{e_1, e_2, \cdots, e_n\}$ is a basis for $\mathbb{R}^n$. This means that (i) the collection is linearly independent, and (ii) every vector in $\mathbb{R}^n$ can be expressed as a linear combination of the collection. We call $\{e_k\}_{k=1}^n$ the standard basis for $\mathbb{R}^n$. Since the basis has n elements, the dimension of $\mathbb{R}^n$ is n.

The collection is also the standard basis for $\mathbb{C}^n$. △

Example 2.2.4 (Another basis for $\mathbb{R}^2$) *The vectors x_1 and x_2 given in Example 2.2.1 form a basis for $\mathbb{R}^2$.*

Example 2.2.5 (Standard basis for $\mathbb{R}^{m\times n}$, $\mathbb{C}^{m\times n}$) *Let M_{ij} be the $m \times n$ matrix whose entries are all zeros except for the (i, j) element which is 1. The collection $\{M_{ij}\}_{i,j=1}^{m,n}$ is the standard basis for $\mathbb{R}^{m\times n}$ and $\mathbb{C}^{m\times n}$.* △

Example 2.2.6 (Standard basis for $\mathcal{P}_n(T)$) *The polynomials $1, t, t^2, \cdots, t^n$ is a basis for the space $\mathcal{P}_n(T)$ of polynomials of degree at most n defined in Example 2.1.6. This is the standard basis for $\mathcal{P}_n(T)$. The dimension of $\mathcal{P}_n(T)$ is $n + 1$.* △

Example 2.2.7 (Bernstein basis for $\mathcal{P}_n([0,1])$) *The Bernstein polynomials of degree n are defined by:*

$$B_i(t) = \frac{n!}{i!(n-i)!} t^i (1-t)^{n-i}$$

for $i = 0, 1, \cdots, n$. These polynomials form a basis for $\mathcal{P}_n([0,1])$. △

Recall the nesting property of span given in Proposition 2.2.2. We use it give the following abstract Gram-Schmidt procedure to compute a basis.

Algorithm 2.2.1 (Abstract Gram-Schmidt) *Let U be a subspace of a vector space V.*

Step 1: *Set iteration counter $k = 1$. Choose a non-zero vector $x_1 \in U$ and set $\mathcal{B}_k = \{x_1\}$.*

Step 2: *If span of $\mathcal{B}_k$ is U, stop. Otherwise, choose a vector x_{k+1} in U but not in the span of $\mathcal{B}_k$.*

Step 3: *Set*

$$\mathcal{B}_{k+1} = \mathcal{B}_k \bigcup \{x_{k+1}\}$$

and increment iteration counter by 1. Go to Step 2. ■

This algorithm terminates if U is finite dimensional. It produces a sequence of sets $\{\mathcal{B}_k\}$ whose spans form a strictly increasing sequence.

2.3 Coordinates and isomorphisms

Elements of vector spaces should be thought of as abstract objects. One of the great consequences of the concept of basis is that it allows us to obtain a concrete representation of a vector as an array of numbers (scalars). We have been using this representation thus far in examples. It turns out that, in advanced engineering problems, the abstract version is essential for problem formulation and the concrete representation is useful in computations. The array representation gives the *coordinates* of the vector in a chosen basis. This concept is made precise below along with the concept of *coordinate change*.

Proposition 2.3.1 (Representation in a basis is unique) *Let $\{x_k\}$ be a basis for a vector space V. Then, for each $x \in V$, there exist a unique set of scalars $\{\alpha_k\}$ such that $x = \sum_k \alpha_k x_k$.* ■

Suppose that $\{\alpha_k\}$ and $\{\beta_k\}$ are two sets of scalars. Define:

$$x = \sum_k \alpha_k x_k \quad \text{and} \quad y = \sum_k \beta_k x_k$$

Then, according to the proposition, $x = y$ if and only if $\alpha_k = \beta_k$ for each k. This means that all the information needed to distinguish the vector x from any other vector is contained in the scalars $\{\alpha_k\}$ (similarly

for y). So, for all practical purposes, the abstraction x can be replaced by the concrete representation $\{\alpha_k\}$. It is customary to stack these scalars one below the other and form a column vector. The vector so obtained is the coordinate of x in the coordinate system defined by $\{x_k\}$.

Definition 2.3.1 (Coordinates and coordinate system) *Let $\{x_k\}$ be a basis for a vector space V and x be an element of V. The unique set of scalars $\{\alpha_k\}$ such that*

$$x = \sum_k \alpha_k x_k$$

is called the coordinate of x in the basis $\{x_k\}$. The basis is called a coordinate system. ⋆

Example 2.3.1 *The standard basis for $\mathbb{R}^2$ is:*

$$\mathcal{B}_s = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$$

and in this basis, let x be represented as:

$$x = 1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Then, the coordinates of x in the standard basis is $\{1, 2\}$ which is usually written as either $(1, 2)$ or as $[1, 2]^T$ in vector form.

Another basis for $\mathbb{R}^2$ is:

$$\mathcal{B}_o = \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$$

In this basis, the same abstract quantity has the representation:

$$x = 1 \begin{bmatrix} 1 \\ 1 \end{bmatrix} + 1 \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

so that the coordinates of x is $\{1, 1\}$ or in the customary form $[1, 1]^T$. △

Thus, coordinates always require a coordinate system to be defined. We usually take the coordinate system to be the one defined by the standard basis and do not explicitly mention it. A change of coordinates is often required to better understand and solve problems more efficiently. In the above example, we considered two bases $\mathcal{B}_s$ and $\mathcal{B}_o$ for $\mathbb{R}^2$. In these bases, the coordinates of the vector x turned out to be $[1, 2]^T$ and $[1, 1]^T$ respectively. That is, when the basis was changed from $\mathcal{B}_s$ to $\mathcal{B}_o$, the coordinates underwent a *transformation*. This transformation is very important and is defined below.

Definition 2.3.2 (Coordinate transformation) *Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be two bases for a vector space V. For $i = 1, 2$, let x_i denote the coordinate of $x \in V$ in the basis $\mathcal{B}_i$. Then, the map*

$$x_1 \mapsto x_2$$

is called the coordinate transformation from $\mathcal{B}_1$ into $\mathcal{B}_2$. ⋆

The definition is shown graphically in Figure 2.2. The map Id_i is an identification map that associates each element of V with an unique element of the span of the elements of the basis $\mathcal{B}_i$. It is an invertible map. We say that the diagram in Figure 2.2 *commutes* because $T \circ Id_1 = Id_2$ where T denotes the coordinate transformation. It gives a way to explicitly compute the coordinate transformation for finite dimensional vector spaces. The next result shows how and summarizes some special cases.

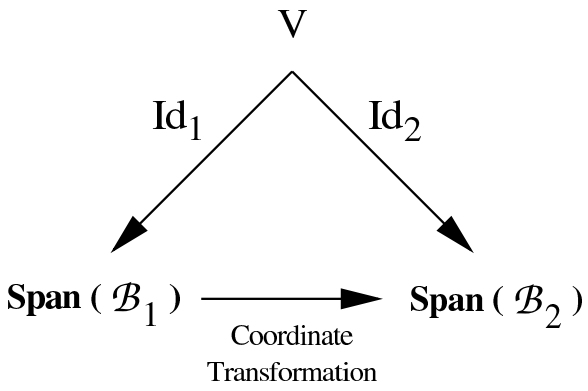

Figure 2.2: Coordinate transformation

Proposition 2.3.2 (Computation of coordinate transformation) *The following statements are true.*

1. *Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be two bases for a vector space V of dimension n. Let T denote the coordinate transformation from the system $\mathcal{B}_1$ into $\mathcal{B}_2$. Then, T is completely defined by its action on the elements of $\mathcal{B}_1$.*

2. *Let*

$$\mathcal{B}_1 = \{x_1, x_2, \cdots, x_n\} \quad \text{and} \quad \mathcal{B}_2 = \{y_1, y_2, \cdots, y_n\}$$

be two bases for $\mathbb{R}^n$. Then, the coordinate transformation T from the system $\mathcal{B}_1$ into $\mathcal{B}_2$ is the solution of:

$$T\,[\,y_1 \quad y_2 \quad \cdots \quad y_n\,] = [\,x_1 \quad x_2 \quad \cdots \quad x_n\,]$$

Moreover, if α is the coordinate of a vector $x \in V$ in the system $\mathcal{B}_1$, then $T\alpha$ is the coordinate of x in the system $\mathcal{B}_2$. ■

We gave several examples of vector spaces including spaces of vectors, matrices and functions. While these spaces are composed of different objects, it is possible to show that under certain conditions they are identical for all purposes. This is evident in Proposition 2.3.1 which allows us to transfer an abstract vector to an array of scalars. It is also used in Proposition 2.3.2.

Definition 2.3.3 (Isomorphism) *Let V and W be vector spaces over the same field. An isomorphism between V and W is a function $f : V \to W$ that is linear, one-to-one and onto.*

When an isomorphism exists between two vector spaces, we say that the vector spaces are isomorphic. ⋆

Theorem 2.3.1 (Isomorphic vector spaces) *Let V and W be vector spaces over the same field. The following statements are equivalent.*

1. *V and W are isomorphic*
2. *The dimension of V is equal to the dimension of W*

Suppose that statement 2 holds. Let $\{v_k\}$ be a basis for V and $\{w_k\}$ be a basis for W. Then the mapping

$$v_k \mapsto w_k \quad \textit{for each } k$$

is an isomorphism between V and W. ■

Accordingly, any real vector space of dimension n is isomorphic to $\mathbb{R}^n$. This is very important as it allows us to convert problems on abstract vector spaces into problems on $\mathbb{R}^n$ and perform calculations on a computer. The following example illustrates this.

Example 2.3.2 (Polynomial manipulations on a computer) *Consider the real vector space $\mathcal{P}_n(T)$ of polynomials of degree at most n defined in Example 2.1.6. Recall that the dimension of $\mathcal{P}_n(T)$ is $n+1$. We will show that $\mathcal{P}_n(T)$ is isomorphic to $\mathbb{R}^{n+1}$ by constructing an isomorphism called an* identification map. *This map will give a straightforward way to manipulate polynomials on a computer.*

To construct an identification map, proceed as follows:

1. *Choose a basis for $\mathcal{P}_n(T)$ and a basis for $\mathbb{R}^{n+1}$, say the standard bases $\{1, t, t^2, \cdots, t^n\}$ for $\mathcal{P}_n(T)$ and $\{e_1, e_2, e_3, \cdots, e_{n+1}\}$ for $\mathbb{R}^{n+1}$. Here, e_k is the column vector of size $n+1$ whose elements are all zeros except for the kth entry which is one.*
2. *Pick a polynomial p in $\mathcal{P}_n(T)$ and express it in the chosen basis:*

$$p(t) = \sum_{k=0}^{n} p_k t^k$$

   where $\{p_0, p_1, \cdots, p_n\}$ is the coordinate of p in the standard basis.

3. *Define the vector in $\mathbb{R}^{n+1}$:*

$$x_p = \sum_{k=1}^{n+1} p_{k-1} e_k = [\, p_0 \quad p_1 \quad \cdots \quad p_n \,]^T$$

   with coordinates $\{p_0, p_1, \cdots, p_n\}$ in the standard basis on $\mathbb{R}^{n+1}$.

4. *Define identification map as the map that takes p to x_p. This map is shown to be linear, one-to-one and onto.*

That is, we stack up the polynomial coefficients to form a vector. Since a polynomial is completely determined by its coefficients, we can uniquely identify a polynomial with its coefficient vector. Thus, to add two polynomials on a computer, we simply write a program that accepts coefficient vectors and returns sum of the input vectors. △

2.4 Convex sets and functions

Convexity is perhaps the most important concept in optimization theory. It enables us to determine global properties by studying local properties. For example, local minimum of a convex function is also its global minimum. Convexity also guarantees the existence of highly efficient algorithms to numerically solve problems. For example, the linear programming problem can be solved efficiently. This section gives a very brief introduction to the subject.

Definition 2.4.1 (Convex set) *Let V be a real vector space. A non-empty subset C of V is convex if and only if whenever x and y are in C, the vector*

$$\alpha x + (1-\alpha)y$$

is in C for any $\alpha \in [0,1]$. ⋆

Given x and y in a real vector space V, the *line* connecting them is the set of all points of the form $\alpha x + (1-\alpha)y$ for $\alpha \in [0,1]$. So, a non-empty set C is convex if and only if the line connecting any two points in C is contained in C.

Example 2.4.1 (Subspaces and affine sets are convex) *Every subspace of a real vector space is a convex set. The affine subsets (see definition in Example 2.1.13) of real vector spaces are convex sets.* △

Example 2.4.2 (A non-convex set) *Any finite set other than $\{0\}$ is not convex. For example, $\{0,1\}$ as a subset of $\mathbb{R}$ is not convex.* △

Example 2.4.3 (Ellipses are convex) *Consider $\mathbb{R}^2$ and the ellipse in $\mathbb{R}^2$:*

$$\mathcal{E} = \left\{ \begin{bmatrix} x \\ y \end{bmatrix} : \frac{(x-x_c)^2}{a^2} + \frac{(y-y_c)^2}{b^2} \leq 1 \right\}$$

where (x_c, y_c) is the center of the ellipse and (a,b) define its major and minor axes.

The shaded region in Figure 2.3 is the ellipse. It is a convex set. △

Example 2.4.4 (Linear equality and inequality) *Let $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$ be given. The set:*

$$\{x \in \mathbb{R}^n : Ax \leq b\},$$

where $\leq$ denotes component-wise less than or equal to, when non-empty is a convex set. △

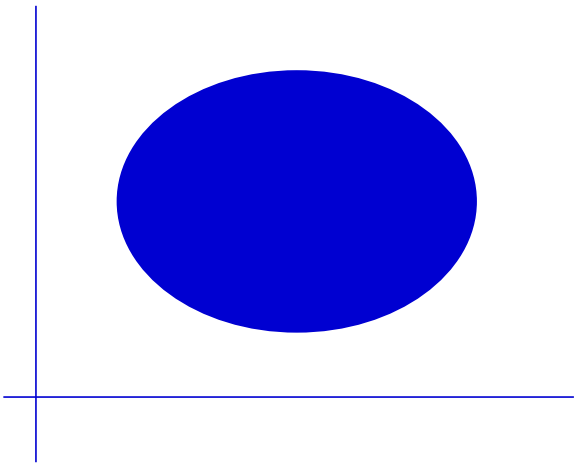

Figure 2.3: An ellipse

Definition 2.4.2 (Convex combination) *Let $\{x_k\}$ be a set of elements in a real vector space V. A convex combination of $\{x_k\}$ is an element x of V that can be expressed as:*

$$x = \sum_{k=1}^{m} \alpha_k x_k \tag{2.5}$$

for some non-negative real numbers $\alpha_1, \alpha_2, \cdots, \alpha_m$ whose sum is 1. ⋆

The main difference between linear and convex combinations is that the scalars are free in the former, while in the latter, they are non-negative and sum to 1. Recall that the set of all linear combinations is the span. We have a similar concept for the set of all convex combinations.

Definition 2.4.3 (Convex hull of elements) *The convex hull of a set of elements $\{x_k\}$ in a real vector space V is*

$$\begin{aligned} \mathbf{Conv}\left(\{x_k\}\right) &= \{x \in V : x \text{ is a convex combination of } \{x_k\}\} && \text{(2.6a)} \\ &= \left\{ \begin{array}{r} x \in V : x = \sum_{k=1}^{m} \alpha_k x_k \text{ for some } \alpha_1 \geq 0, \alpha_2 \geq 0, \cdots, \alpha_m \geq 0 \\ \text{such that } \sum_{k=1}^{m} \alpha_k = 1 \end{array} \right\}, && \text{(2.6b)} \end{aligned}$$

i.e, the set of all elements in V that are convex combinations of $\{x_k\}$ ⋆

Example 2.4.5 *Consider the vector space $\mathbb{R}^2$ and the vectors:*

$$x_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \; x_2 = \begin{bmatrix} 2 \\ 2 \end{bmatrix} \; \text{and} \; x_3 = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

as shown in Figure 2.4. The shaded region in the figure is the convex hull of $\{x_1, x_2, x_3\}$.

The vectors x_1, x_2 and x_3 form the vertices of the triangular region. △

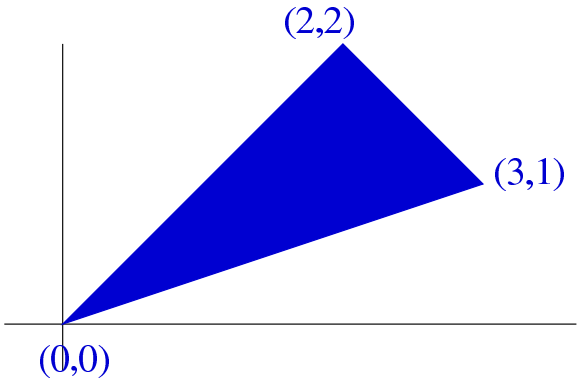

Figure 2.4: An example of convex hull

The shaded region in Figure 2.4 is an example of a polytope. Its definition is given below.

Definition 2.4.4 (Polytope) *Let $\{x_k\}_{k=1}^m$ be a finite set of elements in a real vector space V. The convex hull of $\{x_k\}_{k=1}^m$ is called the polytope generated by $\{x_k\}_{k=1}^m$.* ★

Example 2.4.6 (A polytope of Gaussian density functions) *Consider the real vector space $\mathcal{C}(\mathbb{R})$ of continuous functions on $\mathbb{R}$ taking values in $\mathbb{R}$. For $k = 1, 2, \cdots, m$, define the Gaussian probability density functions:*

$$f_k(x) = \sqrt{\frac{k}{2\pi}} \exp(-\frac{k}{2}x^2) \quad \text{for all } x \in \mathbb{R}$$

The convex hull of $\{f_k\}_{k=1}^m$ is a polytope in $\mathcal{C}(\mathbb{R})$. A typical element of this polytope is a function of the form:

$$\sum_{k=1}^{m} \alpha_k \left(\sqrt{\frac{k}{2\pi}} \exp(-\frac{k}{2}x^2) \right)$$

where the coefficients $\{\alpha_k\}$ are all non-negative and sum to 1. Hence, every element of the polytope is a probability density function. In statistics, elements of this polytope are known as mixtures of Gaussian density functions. △

Another important geometric object is obtained by generalizing the definition of an ellipse (see Example 2.4.3) to arbitrary vector spaces. We give the definition for $\mathbb{R}^n$.

Definition 2.4.5 (Ellipsoid) *Let $P \in \mathbb{R}^{n \times n}$ be a strictly positive matrix and $x_c \in \mathbb{R}^n$. The set*

$$\mathcal{E}(P, x_c) = \left\{ x \in \mathbb{R}^n : (x - x_c)^T P^{-1} (x - x_c) \leq 1 \right\}$$

is called an ellipsoid centered at x_c. It is a convex set. ★

We shall define positive matrices in the next chapter and the above definition will become clearer.

Definition 2.4.6 (Convex function) *Let S be a convex subset of a real vector space V. A function $f : S \to \mathbb{R}$ is said to be convex if*

$$f(\alpha x + (1-\alpha)y) \leq \alpha f(x) + (1-\alpha)f(y)$$

for all $\alpha \in [0,1]$ and x, y in S. ⋆

The definition is shown in Figure 2.5. It says that the function always lies below straight lines that connect any two function values. It is important to note that convexity of a function depends on its domain S. In particular, S must be a convex set.

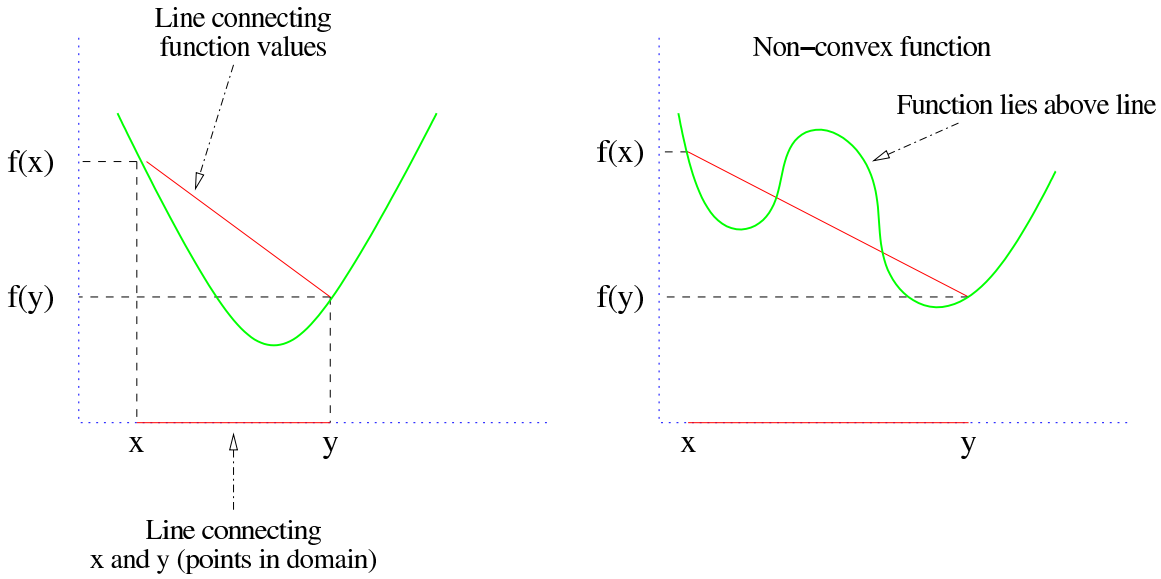

Figure 2.5: Function on the left is convex; function on the right is non-convex

Example 2.4.7 (Linear and affine functions are convex) *Consider $\mathbb{R}^n$. Let a be a vector in $\mathbb{R}^n$ and b be a real number. Define the function $f : \mathbb{R}^n \to \mathbb{R}$ as follows:*

$$f(x) = a^T x + b$$

f is an affine function and when $b = 0$ is a linear function. f is convex.

2.5 Some optimization problems

Let V be a real vector space and S be a subset of V. Let $f : S \to \mathbb{R}$ be a function. Without being precise, the prototypical optimization problem is stated as:

$$\min_{x \in S} f(x)$$

or as

$$\begin{aligned} &\min_{x \in V} f(x) \\ &\text{subject to } x \in S \end{aligned}$$

and consists of two parts:

1. find the minimum value γ_{min}, called the minimum function value, taken by the function f over S, and
2. find a point $x_{\min}$ in S, called the optimal solution, such that $f(x_{\min}) = \gamma_{\min}$, i.e, the function achieves the minimum value at $x_{\min}$.

The notation itself gives the following information. f is the function being minimized, x is the unknown or decision or search variable, and S is a feasible set.

Definition 2.5.1 (Convex optimization problem) *Let S be a convex subset of a real vector space and $f : V \to \mathbb{R}$ be a convex function. The minimization problem:*

$$\begin{aligned} &\min_{x \in V} f(x) \\ &\textit{subject to } x \in S \end{aligned}$$

is called a convex optimization problem. In words, a convex optimization problem is a minimization problem of a convex function over a convex feasible set. ★

If S has a finite dimensional characterization, then convex optimization problem can be efficiently solved. In fact, an algorithm known as elliopsoidal algorithm can be applied universally. We conclude this chapter with some specific examples of convex optimization problems.

Example 2.5.1 (Linear programming problem) *Take $V = \mathbb{R}^n$,*

$$S = \{x : Ax \leq b\}$$

and f to be $f(x) = c^T x$. Here, $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$ are given. The resulting problem:

$$\begin{aligned} &\min_{x \in \mathbb{R}^n} c^T x \\ &\textit{subject to } Ax \leq b \end{aligned}$$

is called the linear programming problem. It can be solved with the ellipsoidal algorithm, simplex method etc. △

Example 2.5.2 (Quadratic programming problem) *Take* $V = \mathbb{R}^n$,

$$S = \{x : Ax \leq b\}$$

and f *to be* $f(x) = x^T Q x + r^T x$. *Here,* $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $Q \in \mathbb{R}^{n \times n}$ *and* $r \in \mathbb{R}^n$ *are given. We assume that* Q *is a strictly positive matrix. The resulting problem:*

$$\begin{array}{l} \min\limits_{x \in \mathbb{R}^n} x^T Q x + r^T x \\ \textit{subject to } Ax \leq b \end{array}$$

is called the quadratic programming problem. △

Chapter 3

Matrix Theory

Matrices arise as *representations* of linear operators and systems. This chapter begins with a review of different types of matrices, and then describes certain numbers and spaces associated with matrices. These numbers and spaces are clear evidence of abstractness of matrices over and beyond their representation as rectangular arrays. Section 3.3 is an introduction to computational linear algebra and is particularly important. The last section presents some applications to linear equations and optimization.

3.1 Different kinds of matrices

We shall refer to matrices in $\mathbb{R}^{m\times n}$ as real matrices and matrices in $\mathbb{C}^{m\times n}$ as complex matrices. The transpose and complex conjugate transpose of a matrix A are denoted by A^T and A^* respectively. If A is a real matrix, then $A^* = A^T$. Some occasions require matrices with infinite number of rows and/or columns and in such cases, we shall explictly say infinite-sized matrix. In all other occasions, matrices will be understood to have finite number of rows and columns.

A matrix $A \in \mathbb{R}^{n\times n}$ with the property $A = A^T$ is called *symmetric*, while a matrix $A \in \mathbb{R}^{n\times n}$ with the property $A = -A^T$ is called *skew-symmetric*. The only matrix that is both symmetric and skew-symmetric is the zero matrix. A matrix $A \in \mathbb{C}^{n\times n}$ with the property $A = A^*$ is called *Hermitian*, while a matrix $A \in \mathbb{C}^{n\times n}$ with the property $A = -A^*$ is called *skew-Hermitian*. Since the complex conjugate of a real matrix is itself, symmetric matrices are Hermitian. It can be shown that the diagonal elements of skew-symmetric and skew-Hermitian matrices are all equal to zero.

A matrix $A \in \mathbb{R}^{n\times n}$ is *orthogonal* if $A^TA = AA^T =$ a diagonal matrix. It is *orthonormal* if $A^TA = AA^T = I$ where I denotes the identity matrix of appropriate dimension. A complex matrix is *inner* if $A^*A = I$ and *co-inner* if $AA^* = I$. A complex matrix is *unitary* if it is inner and co-inner. Inner matrices are tall; co-inner matrices are fat; orthonormal matrices (defined only for real matrices) are unitary.

An *upper-triangular complex Jordan matrix of size n with eigenvalue λ*, denoted by $J_n(\lambda)$, is a square

matrix of the form:

$$J_n(\lambda) = \begin{bmatrix} \lambda & 1 & 0 & \cdot & \cdot & \cdot & 0 & 0 \\ 0 & \lambda & 1 & \cdot & \cdot & \cdot & 0 & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & 0 & \cdot & \cdot & \cdot & \lambda & 1 \\ 0 & 0 & 0 & \cdot & \cdot & \cdot & 0 & \lambda \end{bmatrix}$$

that is, the elements on the main diagonal are all equal to λ, the elements on the diagonal above the main diagonal are all equal to 1 and all the remaining elements are zeros. It is an element of $\mathbb{C}^{n\times n}$. There are lower-triangular and real versions of Jordan matrices, but in these notes, we shall only consider the upper-triangular complex version. So, we simply say Jordan matrix of size n with eigenvalue λ.

A matrix $A \in \mathbb{C}^{n\times n}$ is *invertible* if and only if there exists a matrix $B \in \mathbb{C}^{n\times n}$ such that $AB = I$. In this case, B is called *the inverse of* A and is denoted by A^{-1}. It is easy to see that $AB = I$ if and only if $BA = I$ so that there is nothing special about the ordering of A and B we chose while defining an invertible matrix. Orthogonal and unitary matrices are invertible. In fact, the inverse of an orthogonal (unitary) matrix A is A^T (respectively, A^*).

A symmetric matrix $A \in \mathbb{R}^{n\times n}$ is *positive semi-definite (or simply positive)*, denoted as $A \geq 0$, if and only if

$$x^T Ax \geq 0$$

for all $x \in \mathbb{R}^n$. A symmetric matrix $A \in \mathbb{R}^{n\times n}$ is *positive definite (or strictly positive)*, denoted as $A > 0$, if and only if

$$x^T Ax > 0$$

for all $x \neq 0$. It is very important to note the difference in the definitions. We require strict positivity in the latter definition only for non-zero vectors. A matrix A is *(strictly) negative* if $-A$ is (strictly) positive. We use the notations $A \leq 0$ and $A < 0$ to mean A is negative and strictly negative. Finally, given symmetric matrices A and B, we write $A \geq B$ to mean $A - B \geq 0$ (similarly for $A > B$).

Note that every positive real number has at most two square roots which differ by -1 (a unitary matrix). We can extend this notion of a square root to positive matrices. Let $A \in \mathbb{C}^{n\times n}$ be a positive matrix. Then, there exist a positive integer k and a matrix $B \in \mathbb{C}^{n\times k}$ such that $A = BB^*$ (this can be seen from singular value decomposition (SVD) discussed later). B is called a *square root of* A. Unlike positive numbers, most matrices have an infinite number of square roots. But, like positive numbers, there is a unique positive square root which will be denoted as $A^{1/2}$.

Throughout these notes, we use the word positive to mean greater than or equal to zero, while strictly positive means strictly greater than zero. Similarly for negative and strictly negative.

Proposition 3.1.1 (Some properties of matrices) *Let $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times n}$ and $C \in \mathbb{R}^{n\times k}$. The following statements are true.*

1. *Let A and B be symmetric matrices. Then, $A+B$ and C^TAC are symmetric matrices. Suppose further that A and B commute, that is, $AB = BA$. Then, AB and BA are symmetric matrices.*

2. *Let A and B be orthogonal matrices. Then, AB and BA are orthogonal matrices.*

3. *Let A and B be (strictly) positive matrices. Then, $A+B$ and ABA^T are (strictly) positive matrices, and C^TAC is positive. Moreover, $A+B \geq A$ and $A+B \geq B$.*

4. *There exists a symmetric matrix A (of size strictly greater than 1) such that $A \ngeq 0$, $A \neq 0$ and $A \nleq 0$.*

5. *If A is a real (complex) matrix, then AA^T (respectively, AA^*) is positive.* ■

These properties show some of the actions that do not destroy structure. For example, statement 1 says that addition of symmetric matrices produce only symmetric matrices. This should not be surprising since the set of symmetric matrices is a vector space. Multiplication destroys symmetry in general, but preserves orthogonal-ness.

Statement 4 can be a source of confusion. There is a natural way to order real numbers. For example, given a real number x, we have either $x > 0$ or $0 > x$ or $x = 0$. Thus, we can always tell which side of zero does a given real number lie. Clearly, a real number is a symmetric matrix (of size 1×1). Unfortunately, the ordering relation on real numbers does not extend to the general case of symmetric matrices. We can only define a *partial order* which cannot always tell which "side" of zero does a symmetric matrix lie.

Let us introduce the following notation. Denote by e_{nk} the vector of size $n \times 1$ whose entries are all zeros except the kth entry which is 1. For example,

$$e_{32} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

It should be clear that $\{e_{n1}, e_{n2}, \cdots, e_{nn}\}$ is the standard basis for $\mathbb{C}^n$.

Proposition 3.1.2 (Matrix action) *Let $A \in \mathbb{C}^{m\times n}$. The following statements are true.*

1. *Let the kth column of A be given by $a_k \in \mathbb{C}^m$. Then,*

$$a_k = Ae_{nk}$$

   (that is, A acting on e_{nk} pulls out the kth column).

2. *Let the (i,j) the element of A be given by $a_{ij} \in \mathbb{C}$. Then,*

$$a_{ij} = e_{mi}^* A e_{nj}$$

3. *$A = 0$ if and only if $Ax = 0$ for all $x \in \mathbb{C}^n$.* ■

If we think of a matrix $A \in \mathbb{C}^{m\times n}$ as an array of numbers:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdot & \cdot & \cdot & a_{1n} \\ a_{21} & a_{22} & \cdot & \cdot & \cdot & a_{2n} \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ a_{m1} & a_{m2} & \cdot & \cdot & \cdot & a_{mn} \end{bmatrix},$$

then the matrix is completely defined by the numbers $\{a_{ij}\}$. Statements 1 and 2 show how to extract these numbers by applying A to the elements of standard bases for its input and output spaces. Hence, we say that a matrix is completely defined by its action on elements of a basis for its input space. Statement 3 says that the matrix whose action is to map every vector to the zero vector is the zero matrix.

3.2 Numbers and spaces associated with matrices

There are many numbers that can be attached to a matrix. For example, to each square matrix, we can attach its determinant. Another number of great importance is an eigenvalue. We can also attach spaces with matrices. These numbers and spaces capture properties that are intrinsic to a given matrix.

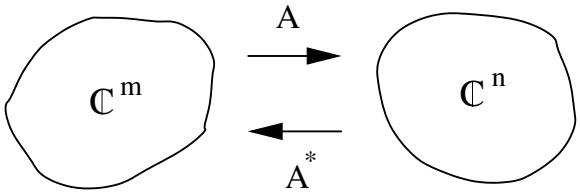

Figure 3.1: Input-output representation of a matrix and its complex conjugate transpose

The input-output representation of a matrix is shown in figure 3.1. It may be useful to write this representation as:

$$y = Ax$$

where $x \in \mathbb{C}^n$ is the input and $y \in \mathbb{C}^m$ is the resulting output. The complex conjugate transpose of $A \in \mathbb{C}^{m\times n}$ is a matrix in $\mathbb{C}^{n\times m}$. Therefore, we can think of A^* as a mapping from $\mathbb{C}^m$ into $\mathbb{C}^n$ as shown in the figure. The input-output representation of A^* may be written as:

$$\tilde{x} = A^*\tilde{y}$$

It is important to note that these representations are of matrices and their complex conjugate transposes, but not of their inverses although at first glance the figure may suggest so.

In what follows, the definitions and results are stated for complex matrices and remain valid for real matrices.

Definition 3.2.1 (Eigenvalue-Eigenvector pair) *Let $A \in \mathbb{C}^{n\times n}$. A pair (λ, x), where $\lambda \in \mathbb{C}$ and $x \in \mathbb{C}^n$, satisfying*

1. $x \neq 0$ *and,*
2. $Ax = \lambda x$

is called an eigenvalue-eigenvector pair of A. λ is an eigenvalue of A and x is an eigenvector associated with λ. ⋆

By definition, an eigenvector is non-zero. This is an important fact that will be used many times. Also, by definition, *every eigenvalue comes with an eigenvector*. But, the eigenvector associated with an eigenvalue is not unique. For example, assume that (λ, x) is an eigenvalue-eigenvector pair for the matrix A. Then, (λ, cx), where c is a non-zero scalar, is also an eigenvalue-eigenvector pair for A (show this by using the definition ?). We say that eigenvectors are *unique up to multiplication by a non-zero scalar*.

Definition 3.2.2 (Characteristic polynomial and equation) *Let $A \in \mathbb{C}^{n\times n}$. The characteristic polynomial of A is the polynomial $p_A : \mathbb{C} \to \mathbb{C}$ defined as:*

$$p_A(s) = \det(sI - A) \quad \text{for all } s \in \mathbb{C}$$

It has degree n. The characteristic equation of A is $p_A(s) = 0$.

An important fact is that the roots of the characteristic polynomial of A are the eigenvalues of A. So, an $n \times n$ matrix can have at most n eigenvalues (this is a consequence of the fundamental theorem of algebra, later we shall give a geometric argument). It is generally difficult to use characteristic equation to compute eigenvalues. We must appeal to the definition for ease; in fact, most numerical methods do.

Example 3.2.1 (Eigenvalues of Hermitian and skew-Hermitian matrices) *Let A be Hermitian. The eigenvalues of A are real numbers. To see this, pick an eigenvalue-eigenvector pair (λ, x) of A. Then, $x \neq 0$, $Ax = \lambda x$ and:*

$$\begin{aligned} \lambda x^* x &= x^* Ax = x^* A^* x = (Ax)^* x \\ &= (\lambda x)^* x = \overline{\lambda} x^* x \end{aligned}$$

where we used the fact that $A = A^*$. The above equation implies:

$$\left(\lambda - \overline{\lambda}\right) x^* x = 0$$

We claim that $x^*x \neq 0$. This can be seen by writing x in terms of its components and carrying out the multiplication x^*x and using the fact that $x \neq 0$. Therefore,

$$\left(\lambda - \overline{\lambda}\right) x^* x = 0$$

implies that

$$\left(\lambda - \overline{\lambda}\right) = 0$$

which means that λ is a real number.

Suppose that M is skew-Hermitian. Then, using similar arguments, we can show that the eigenvalues of M have zero real parts (purely imaginary). $\triangle$

Consider the matrices

$$A_1 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad A_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \tag{3.1}$$

which have the same set of eigenvalues, namely 0 repeated twice. A_1 has two linearly independent eigenvectors; but A_2 has only one eigenvector (up to multiplication by a non-zero scalar). This observation is tied to the concept of multiplicity.

Definition 3.2.3 (Algebraic and geometric multiplicities) *Let λ be an eigenvalue of $A \in \mathbb{C}^{n\times n}$. The algebraic multiplicity of λ, denoted by $\mathrm{alg}_A(\lambda)$, is the number of times it repeats as a root of the characteristic polynomial of A. The geometric multiplicity of λ, denoted by $\mathrm{geo}_A(\lambda)$, is the number of linearly independent eigenvectors associated with λ.* $\star$

Definition 3.2.4 (Defective and non-defective matrices) *Let $A \in \mathbb{C}^{n\times n}$. We say that A is defective if and only if the algebraic and geometric multiplicities of some eigenvalue of A are not equal. A matrix that is not defective is called non-defective (that is, for each eigenvalue of A, its algebraic and geometric multiplicities are equal).* $\star$

Recall that an eigenvalue comes with an eigenvector by definition. So, the number of times an eigenvalue repeats must be at least as large as the number of linearly independent eigenvectors, but it could be strictly larger. In other words,

$$\mathrm{alg}_A(\lambda) \geq \mathrm{geo}_A(\lambda) \quad \text{always.}$$

In the case of a defective matrix, there is at least one eigenvalue for which the above inequality is strict. On the other hand, the inequality is an equality for all the eigenvalues of a non-defective matrix.

Example 3.2.2 *The matrices A_1 and A_2 given in (3.1) have:*

$$\mathrm{geo}_{A_1}(0) = 2 \ \textit{and}\ \mathrm{geo}_{A_2}(0) = 1$$

and

$$\mathrm{alg}_{A_1}(0) = 2 \ \textit{and}\ \mathrm{alg}_{A_2}(0) = 2$$

Hence, A_1 is non-defective and A_2 is defective. Note that A_2 is the jordan matrix $J_2(0)$ of size 2 with eigenvalue 0. It can be shown that $J_n(\lambda)$ with $n > 1$ is defective. $\triangle$

Proposition 3.2.1 (Some non-defective matrices) *Hermitian, skew-Hermitian and unitary matrices are non-defective.* $\blacksquare$

Defective matrices are very difficult to handle computationally. This is because there is a non-defective matrix that lies arbitrarily close to a given defective matrix. For example,

$$\widehat{A}_2 = \begin{bmatrix} \epsilon & 1 \\ 0 & 0 \end{bmatrix},$$

where ϵ is a small non-zero number, is non-defective and lies close to A_2 given in (3.1). Thus, it is virtually impossible to reliably determine multiplicities of defective matrices on a digital computer. The matrix-types listed in the above proposition are non-defective and, hence, computer-friendly.

Definition 3.2.5 (Schmidt pairs and singular values) *Let $A \in \mathbb{C}^{m\times n}$. A singular value σ of A is a positive number for which there exists a pair (u, v) of non-zero vectors, where $u \in \mathbb{C}^m$ and $v \in \mathbb{C}^n$, satisfying*

$$A^*u = \sigma v \quad \textit{and} \quad Av = \sigma u \tag{3.2}$$

The pair (u, v) is called a Schmidt pair associated with σ. ⋆

Let us make the following observation. Suppose that (u, v) is a Schmidt pair of A. Then, by definition, (3.2) holds and

$$(AA^*)u = A(A^*u) = A(\sigma v) = \sigma Av = \sigma^2 u$$

which implies, since u is non-zero, that (σ^2, u) is an eigenvalue-eigenvector pair of AA^*. Similarly, we can show that (σ^2, v) is an eigenvalue-eigenvector pair of A^*A. Thus, *$\sigma \geq 0$ is a singular value of A if and only if σ^2 is an eigenvalue of AA^* (and A^*A).* Note also that Schmidt pairs are eigenvectors of AA^* and A^*A.

Example 3.2.3 *A matrix whose eigenvalues are all zeros need not be zero. But, a matrix is zero if and only if its largest singular value is zero.* △

Definition 3.2.6 (Range space and null space of a matrix) *Let $A \in \mathbb{C}^{m\times n}$. The range space (or simply range) of A, denoted by $\mathcal{R}(A)$, is the set of all elements in $\mathbb{C}^m$ that can be reached from $\mathbb{C}^n$ by applying A, that is,*

$$\begin{aligned} \mathcal{R}(A) &\triangleq \{y \in \mathbb{C}^m : \textit{ there exists } x \in \mathbb{C}^n \textit{ such that } y = Ax\} \\ &\triangleq \{Ax : x \in \mathbb{C}^n\} \end{aligned}$$

The null space of A, denoted by $\mathcal{N}(A)$, is given by:

$$\mathcal{N}(A) \triangleq \{x \in \mathbb{C}^n : Ax = 0\}$$

that is, the set of all elements in $\mathbb{C}^n$ that are mapped to zero by A. ⋆

Range and null spaces are very important. Many engineering problems reduce to computing these spaces efficiently. Examples include data compression and separating signal from noise. Note that the definitions

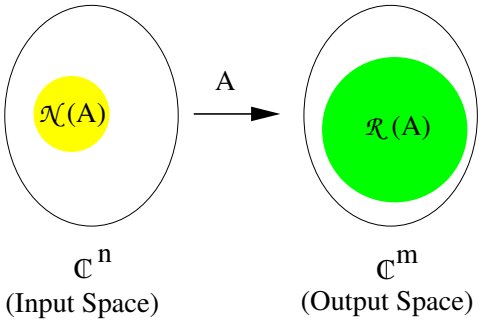

Figure 3.2: Location of range and null space of A

involve systems of linear equations, namely $Ax = y$ in the case of range and $Ax = 0$ in the case of null space. So, these spaces are fundamental objects in the theory of linear equations and more generally of optimization theory.

Let us denote the kth column of A by a_k and the kth element of $x \in \mathbb{C}^n$ by x_k. So,

$$A = [\, a_1 \quad a_2 \quad \cdots \quad a_n \,] \text{ and } x = [\, x_1 \quad x_2 \quad \cdots \quad x_n \,]^T$$

and

$$Ax = a_1x_1 + a_2x_2 + \cdots + a_nx_n = x_1a_1 + x_2a_2 + \cdots + x_na_n$$

where the quantity on the extreme right is the weighted sum of the vectors $\{a_k\}$, in other words, a linear combination of the columns of A. Now, range of A is the set of all vectors that can written as Ax as x ranges over $\mathbb{C}^n$. This leads us to conclude that the range of A is the set of all linear combinations of the columns of A, i.e. the span of the columns of A. For this reason, *range of a matrix is also known as the column span*.

The definition of null space sports the equation $Ax = 0$. In the notation, we have just introduced, this equation becomes:

$$x_1a_1 + x_2a_2 + \cdots + x_na_n = 0$$

Recall that an equation of this form appeared in the definition of linear independence in Chapter 2. In fact, if the only solution of the above equation is $x_1 = x_2 = \cdots = 0$, then the columns of A are linearly independent, while a non-zero solution (meaning at least one of x_k's is non-zero) implies that the columns are linearly dependent. Thus, the null space of A tells us if the columns of A are linearly independent or not.

Example 3.2.4 *Consider the matrices A_1 and A_2 given in (3.1). We have:*

$$\mathcal{R}(A_1) \triangleq \left\{ \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} x : x \in \mathbb{C}^2 \right\} = \{0\}$$

and

$$\mathcal{R}(A_2) \triangleq \left\{ \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x : x \in \mathbb{C}^2 \right\}$$

$$= \left\{ \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} : x_1 \in \mathbb{C}, x_2 \in \mathbb{C} \right\}$$
$$= \left\{ \begin{bmatrix} x_2 \\ 0 \end{bmatrix} : x_2 \in \mathbb{C} \right\} = \mathbf{Span}\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right)$$

(recall the definition of span from (2.4)). Similar computations yield the null spaces. △

We already mentioned that the range of a matrix is the span of its columns. Since span of a collection of vectors is a subspace, range of a matrix is a subspace. Now, consider the null space of a matrix A. Take any pair of elements x, y in $\mathcal{N}(A)$. Then, by definition,

$$Ax = 0 \text{ and } Ay = 0$$

So, for any pair of scalars α and β, we have

$$A(\alpha x + \beta y) = \alpha Ax + \beta Ay = 0$$

which proves that $\alpha x + \beta y$ is in the null space. Therefore, the null space of A is also a subspace.

Theorem 3.2.1 (Range and null space are subspaces) *Let $A \in \mathbb{C}^{m \times n}$. The range of A is a subspace of $\mathbb{C}^m$ and the null space of A is a subspace of $\mathbb{C}^n$.* ■

It is important to note that the range and null space of A are in general subspaces of different spaces (respectively, output space and input space). See figure 3.2. A great simplification occurs in the case of a square matrix $A \in \mathbb{C}^{n \times n}$. In this case, $\mathbb{C}^n$ is both the input space and the output space. Although inputs and outputs may correspond to quantities that are physically different, by isomorphism, the spaces are really two copies of the same space.

Definition 3.2.7 (Rank and nullity of a matrix) *Let $A \in \mathbb{C}^{m \times n}$. The rank of A is the dimension of the range of A. The nullity of A is the dimension of the null space of A.* ⋆

If A has n columns, then the span of the columns of A cannot have dimension greater than n. So, the rank of a matrix is less than or equal to the number of columns. Rank is also less than or equal to the number of rows, but we have to wait a while to see it.

Proposition 3.2.2 *Let $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{n \times k}$. The range of AB is contained in the range of A. The null space of AB contains the null space of B.* ■

We conclude this section with another number associated with a square matrix.

Definition 3.2.8 (Trace of a square matrix) *Let $A \in \mathbb{C}^{n \times n}$. The trace of A, denoted by $\mathrm{Tr}(A)$, is the sum of the diagonal elements of A.* ⋆

Trace and determinants are related to the other numbers introduced earlier.

Proposition 3.2.3 (Properties of determinant and trace) *Let $A \in \mathbb{C}^{n\times n}$ and $B \in \mathbb{C}^{n\times n}$. The following statements are true.*

1. *The determinant of A is the product of the eigenvalues of A. The trace of A is the sum of the eigenvalues of A.*
2. *The determinant of AB is equal to the product of the determinants of A and B.*
3. $\mathrm{Tr}\,(AB) = \mathrm{Tr}\left(B^T A^T\right) = \mathrm{Tr}\,(BA)$ ■

3.3 Decompositions

Decompositions are used to reveal properties of matrices. An easy decomposition that follows from the vector space structure of the set of real square matrices is the additive decomposition of a matrix into a symmetric part and a skew-symmetric part. That is, $A \in \mathbb{R}^{n\times n}$ can be written as:

$$A = X + Y$$

where X is symmetric and Y is skew-symmetric. In fact,

$$X = \frac{1}{2}\left(A + A^T\right) \text{ and } Y = \frac{1}{2}\left(A - A^T\right)$$

We shall examine decompositions that are multiplicative, that is, those that express a given matrix as the product of other matrices. The most important such decompositions are QR, Jordan (or spectral), singular value, and Schur. The Jordan decomposition for non-defective matrices becomes the eigenvalue-eigenvector decomposition. QR and Schur decompositions are the easiest to compute (on a digital computer) followed by singular value decomposition (SVD). Jordan decompositions are next to impossible to compute, but are very useful in proving theorems.

3.3.1 QR Factorization

Consider a system of linear equations of the form $Ax = b$ where x is the unknown variable. In many situations, A has more rows than columns meaning that there are more equations to be satisfied than the number of unknowns. QR factorization is useful in solving the system of equations in such cases.

Theorem 3.3.1 (QR factorization) *Let $A \in \mathbb{C}^{m\times n}$ with $m \geq n$. There exist a positive integer $r \leq n$, an unitary matrix $Q \in \mathbb{C}^{m\times m}$, a permutation matrix $\Pi \in \mathbb{C}^{n\times n}$ and a matrix $R \in \mathbb{C}^{m\times n}$ with the structure*

$$R = \begin{bmatrix} R_{11} & R_{12} \\ 0_{(m-r)\times r} & 0_{(m-r)\times(n-r)} \end{bmatrix},$$

where R_{11} is upper-triangular and has rank r, such that $A\Pi = QR$. ■

The matrices Q and R are called the QR factors of A, and r is the rank of A. There are many efficient procedures for computing QR [6]. The most important uses of QR factorization are in computing the range and null spaces of matrices and in solving linear equations.

Theorem 3.3.2 (Properties of QR) *Let $A \in \mathbb{C}^{m\times n}$ with $m \geq n$ and (Q, R) be the QR factors of A as in Theorem 3.3.1. Let r be the rank of A and partition Q as:*

$$Q = [\, Q_1 \quad Q_2 \,]$$

where Q_1 consists of the first r columns of Q and Q_2 consists of the remaining $m - r$ columns of Q. The following statements are true.

1. *A basis for the range of A is given by the columns of Q_1. That is, $\mathcal{R}(A) = \mathcal{R}(Q_1)$*
2. *A basis for the null space of A^* is given by the columns of Q_2. That is, $\mathcal{N}(A^*) = \mathcal{R}(Q_2)$*
3. *$A = Q_1 [\, R_{11} \quad R_{12} \,]$ where R_{11} and R_{12} are as in Theorem 3.3.1.* ■

The basis given by QR consists of vectors of unit length and mutually orthogonal (we have not yet defined orthogonality). This follows from the fact that Q is unitary.

3.3.2 Singular Value Decomposition (SVD)

Schmidt pairs lead to the following decomposition known as Singular Value Decomposition (SVD). It is perhaps the most important decomposition and is frequently used in data compression, document searching etc.

Theorem 3.3.3 (Singular Value Decomposition (SVD)) *Let $A \in \mathbb{C}^{m\times n}$. There exist a positive integer $r \leq \min\{m, n\}$ and unitary matrices $U \in \mathbb{C}^{m\times m}$ and $V \in \mathbb{C}^{n\times n}$ such that*

$$A = U \begin{bmatrix} \Sigma_1 & 0_{r\times(n-r)} \\ 0_{(m-r)\times r} & 0_{(m-r)\times(n-r)} \end{bmatrix} V^* \tag{3.3}$$

where Σ_1 is an $r \times r$ diagonal matrix:

$$\Sigma_1 = \begin{bmatrix} \sigma_1 & 0 & 0 & \cdot & \cdot & \cdot & 0 \\ 0 & \sigma_2 & 0 & \cdot & \cdot & \cdot & 0 \\ 0 & 0 & \sigma_3 & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 0 \\ 0 & 0 & 0 & \cdot & \cdot & \cdot & \sigma_r \end{bmatrix} \tag{3.4}$$

whose diagonal elements satisfy $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$ (that is, they are all strictly positive and are ordered from the largest to the smallest). ■

σ_k is called the kth non-zero singular value of A. Singular values were defined earlier and are the positive square roots of the eigenvalues of AA^*. By our convention, σ_1 denotes the largest singular value and σ_r is the smallest non-zero (when one exists) singular value. The total number of strictly positive singular values is r and it is equal to the rank of A.

Theorem 3.3.4 (Properties of SVD) *Let $A \in \mathbb{C}^{m\times n}$ and consider its SVD given in Theorem 3.3.3. Partition U and V as:*

$$U = [\, U_1 \quad U_2 \,] \text{ and } V = [\, V_1 \quad V_2 \,]$$

where U_1 and V_1 consist of the first r columns of U and V respectively, and U_2 and V_2 consist of the remaining $m-r$ columns of U and the remaining $n-r$ columns of V respectively. The following statements are true.

1. *A basis for the range of A is given by the columns of U_1. That is, $\mathcal{R}(A) = \mathcal{R}(U_1)$*
2. *A basis for the null space of A is given by the columns of V_2. That is, $\mathcal{N}(A) = \mathcal{R}(V_2)$*
3. *A basis for the range of A^* is given by the columns of V_1. That is, $\mathcal{R}(A^*) = \mathcal{R}(V_1)$*
4. *A basis for the null space of A^* is given by the columns of U_2. That is, $\mathcal{N}(A^*) = \mathcal{R}(U_2)$*
5. *The columns of U and V are eigenvectors of AA^* and A^*A respectively. That is,*

$$AA^*U = U \begin{bmatrix} \Sigma_1^2 & 0_{r\times(m-r)} \\ 0_{(m-r)\times r} & 0_{(m-r)\times(m-r)} \end{bmatrix}$$

and

$$A^*AV = V \begin{bmatrix} \Sigma_1^2 & 0_{r\times(n-r)} \\ 0_{(n-r)\times r} & 0_{(n-r)\times(n-r)} \end{bmatrix}$$

6. *The matrix A can be written as:*

$$A = U_1 \Sigma_1 V_1^*$$

where Σ_1 is as in Theorem 3.3.1. ■

3.3.3 Spectral Decomposition

The *spectrum* of a matrix is the set of its eigenvalues. Spectral decomposition is a decomposition that explicitly shows the eigenvalues and their multiplicity structures.

Theorem 3.3.5 (Complex Jordan form) *Let $A \in \mathbb{C}^{n\times n}$. There exists an invertible matrix M such that*

$$A = MJM^{-1}$$

where

$$J = \begin{bmatrix} J_{n_1}(\lambda_1) & 0 & \cdots & 0 \\ 0 & J_{n_2}(\lambda_2) & \cdots & 0 \\ \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ 0 & 0 & \cdots & J_{n_m}(\lambda_m) \end{bmatrix}$$

and $J_{n_k}(\lambda_k)$ is the $n_k \times n_k$ upper-triangular Jordan matrix with eigenvalue λ_k. ■

The block-diagonal matrix J is called the Jordan form of A. It is a block-upper-triangular matrix. So, the λ_k's are eigenvalues of A. Every matrix can be put in its Jordan form. To put a matrix in its Jordan form, we compute a specific similarity transform and apply the transformation. This may be a difficult task because, in general, matrices can be defective. But, the Jordan forms of non-defective matrices can be computed more reliably and turn out to be diagonal. It is often called the *eigenvalue-eigenvector decomposition*.

Theorem 3.3.6 (Eigenvalue-eigenvector decomposition) *Let $A \in \mathbb{C}^{n\times n}$ be non-defective. There exists an invertible matrix E such that*

$$A = E\Lambda E^{-1}$$

where

$$\Lambda = \text{diag } [\lambda_1, \lambda_2, \cdots, \lambda_n]$$

and λ_k's are eigenvalues of A. ■

We can write the eigenvalue-eigenvector decomposition in the more revealing form:

$$AE = E\Lambda$$

Now, let v_k denote the kth column of E:

$$E = [\, v_1 \quad v_2 \quad \cdots \quad v_n \,]$$

and carry out the matrix multiplications using the diagonal structure of Λ to obtain:

$$AE = [\, Av_1 \quad Av_2 \quad \cdots \quad Av_n \,] \quad \text{and} \quad E\Lambda = [\, \lambda_1 v_1 \quad \lambda_2 v_2 \quad \cdots \quad \lambda_n v_n \,]$$

so that the eigenvalue-eigenvector decomposition is really

$$Av_1 = \lambda_1 v_1,\ Av_2 = \lambda_2 v_2,\ \cdots,\ Av_n = \lambda_n v_n$$

This shows that the columns of E are the eigenvectors associated with the eigenvalues of A. Hence, the terminology eigenvalue-eigenvector decomposition. Similar manipulations when carried out with the Jordan decomposition lead to the notion of a *generalized eigenvector*.

Example 3.3.1 (Hermitian matrices) *Let $A \in \mathbb{C}^{n\times n}$ be Hermitian. Then, the eigenvalue-eigenvector decomposition of A has the form:*

$$A = E\Lambda E^*$$

where the diagonal elements of Λ are real numbers. △

3.3.4 Similarity

Note that in the Jordan decomposition (or eigenvalue-eigenvactor decomposition), we begin with a matrix A and end up with a matrix J (respectively Λ) that reveals some intrinsic structure of A. In fact, this procedure is a special case of what is known as similarity transformation.

Definition 3.3.1 (Similarity transformation, similar matrices) *Let $M \in \mathbb{C}^{n\times n}$ be an invertible matrix. The map that takes a matrix $A \in \mathbb{C}^{n\times n}$ into MAM^{-1} is called a similarity transformation.*

Two matrices $A \in \mathbb{C}^{n\times n}$ and $\hat{A} \in \mathbb{C}^{n\times n}$ are said to be similar if and only if there is a similarity transformation M such that $\hat{A} = MAM^{-1}$. ⋆

Example 3.3.2 (Eigenvalue-Eigenvector decomposition) *Let*

$$A = E\Lambda E^{-1}$$

be the eigenvalue-eigenvector decomposition of A. Then, A and Λ are similar, and E is the similarity transform. △

Proposition 3.3.1 (Eigen-structures are similarity-invariant) *Suppose that $A \in \mathbb{C}^{n\times n}$ and $\hat{A} \in \mathbb{C}^{n\times n}$ are similar. Then, the eigenvalues, their geometric and algebraic multiplicities of A and $\hat{A}$ are the same.* ■

Example 3.3.3 (Hermitian matrices) *Let $A \in \mathbb{C}^{n\times n}$ and $\hat{A} \in \mathbb{C}^{n\times n}$ be Hermitian matrices. Then, A and $\hat{A}$ are similar if and only if their eigenvalues are equal counting multiplicities. To see this, suppose that the eigenvalues are same. Apply the eigenvalue-eigenvector decomposition to A and $\hat{A}$:*

$$A = E\Lambda E^{-1} \text{ and } \hat{A} = \hat{E}\Lambda\hat{E}^{-1}$$

where the diagonal matrix of eigenvalues are the same, but the similarity transforms E and $\hat{E}$ could be different. Solving for Λ from the first equation, we get

$$\Lambda = E^{-1}AE$$

Substituting for Λ into the second equation gives

$$\hat{A} = \hat{E}E^{-1}AE\hat{E}^{-1} = \left(\hat{E}E^{-1}\right)A\left(E\hat{E}^{-1}\right)$$

Now, define

$$M = \hat{E}E^{-1}$$

and note that it is invertible. So, A and $\hat{A}$ are similar. The reverse implication is easy and follows from Proposition 3.3.1. △

The matrices A_1 and A_2 given in (3.1) have same eigenvalues, but different Jordan structure. So, A_1 and A_2 are not similar.

3.4 Linear equations and optimization problems

We present a complete analysis of system of linear equations in this section. Although the presentation considers equations involving matrices, it can be easily extended to the more general case of linear operators. Consider the following linear system of equations:

$$Ax = b \tag{3.5}$$

where $A \in \mathbb{C}^{m \times n}$. We shall ask four questions:

1. Does (3.5) have a solution for a given b ?
2. Does (3.5) have a solution for any b ?
3. If a solution exists for a specific b, then is it unique ?
4. If a solution exists for a specific b, then what is a characterization of all solutions (meaning find all solutions) ?

The answer to the first question is the following. A solution exists if and only if b is in the range of A. Recall that the range of A is the set of all points in the output space that can be reached from the input space through the application of A. So, in order to have a solution (put it differently, in order to reach b from the input space), b must be in the range of A. This clearly indicates the answer to the second question which asks if we can reach every point in the output space. A solution exists for any $b \in \mathbb{C}^m$ if and only if the range of $A = \mathbb{C}^m$, i.e. range of A is equal to the entire output space. The third question can be answered using linearity. All solutions can be written as the sum of a particular solution and elements of the null space of A. This is because points in the null space of A "contribute nothing towards b". Note that if the null space contains non-zero elements, then the solution is not unique. So, when a solution exists, it is unique if and only if the null space is $\{0\}$. We summarize these in the following theorem:

Theorem 3.4.1 (Linear equations - existence and uniqueness) *Let $A \in \mathbb{C}^{m \times n}$. The following statements are true.*

1. *Let $b \in \mathbb{C}^m$ be given. The system (3.5) is solvable if and only if $b \in \mathcal{R}(A)$.*
2. *The system (3.5) is solvable for any $b \in \mathbb{C}^m$ if and only if $\mathcal{R}(A) = \mathbb{C}^m$.*
3. *Let $b \in \mathcal{R}(A)$. The system (3.5) has a unique solution if and only if $\mathcal{N}(A) = \{0\}$. If the null space is not equal to $\{0\}$, then all solutions of (3.5) are of the form:*

$$x = x_p + v$$

where v is in $\mathcal{N}(A)$ and $x_p \in \mathbb{C}^n$, called the particular solution, satisfies $Ax_p = b$. ■

This theorem applies to a variety of situations including linear operators. We had shown earlier how to compute a basis for the range and null space of a matrix using SVD (Theorem 3.3.4). So, the tests stated in the theorem can be conducted very efficiently. Our next objective is to state a computational version of the theorem.

Definition 3.4.1 (Pseudo inverse and Moore-Penrose inverse) *Let $A \in \mathbb{C}^{m\times n}$. A pseudo-inverse of A is a matrix $B \in \mathbb{C}^{n\times m}$ satisfying*

$$ABA = A \quad \textit{and}\ BAB = B \tag{3.6}$$

A pseudo-inverse B satisfying

$$(AB)^* = AB \quad \textit{and}\ (BA)^* = BA \tag{3.7}$$

is called the Moore-Penrose inverse of A and is denoted by A^+. ⋆

The Moore-Penrose inverse of a matrix always exists and is unique; but pseudo-inverse is not necessarily unique. When the matrix A is square and invertible, the pseudo- and Moore-Penrose inverses reduce to the inverse A^{-1}. This is easily seen by taking $B = A^{-1}$ and verifying that the conditions in (3.6-3.7) hold.

Proposition 3.4.1 (Formula for Moore-Penrose inverse) *Let the SVD of $A \in \mathbb{C}^{m\times n}$ be as in Theorem 3.3.3. Partition the unitary matrices U and V as in Theorem 3.3.4. The Moore-Penrose inverse of A is given by:*

$$A^+ = V_1 \Sigma_1^{-1} U_1^*$$

where Σ_1, appearing in (3.4), is the $r \times r$ diagonal matrix whose diagonal entries are the singular values of A. ■

We are now ready to summarize the existence theorem 3.4.1 into a computational scheme for solving the linear system of equations.

Theorem 3.4.2 (Solver for linear equations) *Let $A \in \mathbb{C}^{m\times n}$ and $b \in \mathbb{C}^m$ be given. The following statements are equivalent.*

1. *There exists $x \in \mathbb{C}^n$ such that $Ax = b$.*
2. *$AA^+b = b$*

Moreover, if statement 2 holds, then all solutions are of the form:

$$\begin{aligned} x &= A^+b + \left(I - A^+A\right)v && \text{(3.8a)}\\ &= A^+b + V_2V_2^*v && \text{(3.8b)}\\ &= A^+b + V_2z && \text{(3.8c)} \end{aligned}$$

where v and z are arbitrary vectors and A^+ is the Moore-Penrose inverse of A. Further, the solution is unique if $A^+A = I$. ■

Thus, to solve a linear system, we check if statement 2 holds. If so, then there is at least one solution and all solutions are given by the formula (3.8). On the other hand, if statement 2 does not hold, then no solution exists. As given in the theorem, an SVD of A is needed to compute the Moore-Penrose inverse and check statement 2. But, using the formula for A^+ given in Proposition 3.4.1 and the formula for A given in statement 7 of Theorem 3.3.4, we can simplify the existence test to

$$U_1 U_1^* b = b$$

which does not involve computing the Moore-Penrose inverse explicitly. Similarly, the formula for all solutions (3.8) can also be simplified as:

$$x = V_1 \Sigma_1^{-1} U_1^* b + V_2 V_2^* v$$

which involves fewer computations. When A is an arbitrary matrix, this is perhaps the best way (in terms of computations and numerical stability) to solve a linear system. If A has some structure such as Hermitian, then faster methods are available [6]. The solution forms given by (3.8a) and (3.8b) represent each solution as an orthogonal sum and is useful in solving certain optimization problems. The form given by (3.8c) is also useful in that it parameterizes the set solutions with the minimum number of free variables.

Every matrix has a Moore-Penrose inverse. So, we can perform the calculations indicated in the formula (3.8) for all solutions irrespective of whether the linear system has a solution or not. When the linear system has no solution, what is the meaning of x given by the formula (3.8) ? On the other hand, if the linear system has many solutions, which solution should we choose ?

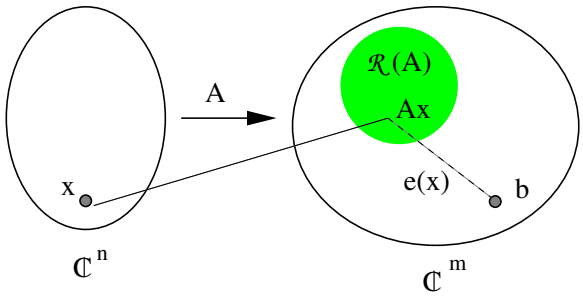

Figure 3.3: Minimization problem

To answer the above questions, let us approach the formula (3.8) from another point of view. Pick any $x \in \mathbb{C}^n$. Then, the quantity

$$e(x) = b - Ax$$

is the error commited by choosing x as a candidate for solution. The quantity

$$e(x)^* e(x) = (b - Ax)^* (b - Ax)$$

is the square error associated with choosing x as the solution. Typically, we would like to find an x that minimizes this error. The situation is shown in Figure 3.3 and can be stated formally as:

$$\min_{x \in \mathbb{C}^n} (b - Ax)^* (b - Ax) \tag{3.9}$$

The square error function can be expanded as follows:

$$(b - Ax)^* (b - Ax) = b^*b + x^*(A^*A)x - (2b^*A)x$$

where b^*b is a constant (independent of the unknown variable x). Therefore, the minimization problem is equivalent to:

$$\min_{x \in \mathbb{C}^n} x^*Qx + r^*x \tag{3.10}$$

with $Q = A^*A$ and $r = -2A^*b$. This is an example of an *unconstrained convex quadratic optimization problem* given in Chapter 2. Now, if there is a solution to the linear system $Ax = b$, then the optimal solution of the minimization problem will also solve $Ax = b$.

Suppose that there is a solution to the linear system $Ax = b$. Then, the formula (3.8) gives all the solutions. In this case, we might ask for the solution of smallest size. That is,

$$\begin{aligned} &\min_{x \in \mathbb{C}^n} x^*x \\ &\text{subject to } Ax = b \end{aligned} \tag{3.11}$$

which is an example of an *equality-constrained convex quadratic optimization problem* (with $Q = I$ and $r = 0$). The motivation for this problem comes from the fact that, in engineering, the decision variable x corresponds to physical variables which should be kept small.

Finally, we could ask for the smallest sized x that minimizes the error $b - Ax$. Here, the size of x is given by x^*x. The solutions to all these problems are stated below.

Theorem 3.4.3 (Solution of optimization problems) *Let $A \in \mathbb{C}^{m \times n}$ and $b \in \mathbb{C}^m$ be given. The following statements are true.*

1. *Any x of the form (3.8) is a solution of the unconstrained optimization problem (3.9). The optimal cost is given by*
$$b^* \left(I - AA^+\right) b$$
2. *The constrained optimization problem (3.11) has a solution if and only if $AA^+b = b$. If it has a solution, then the optimal solution is unique and it is given by $x = A^+b$.*
3. *$x = A^+b$ is the smallest-sized x that minimizes the square error $(b - Ax)^* (b - Ax)$.* ■

3.5 A collection of important results

We conclude this chapter with a number of results that are frequently used.

Theorem 3.5.1 (Properties of range and null space) *Let $A \in \mathbb{C}^{m \times n}$. The following statement are true.*

1. *The direct sum of the range of A and the null space of A^* is $\mathbb{C}^m$, i.e,*

$$\mathcal{R}(A) \oplus \mathcal{N}(A^*) = \mathbb{C}^m$$

Moreover, the sum of the dimensions of $\mathcal{R}(A)$ and $\mathcal{N}(A^*)$ is equal to m.

2. *The direct sum of the range of A^* and the null space of A is $\mathbb{C}^n$, i.e,*

$$\mathcal{R}(A^*) \oplus \mathcal{N}(A) = \mathbb{C}^n$$

Moreover, the sum of the dimensions of $\mathcal{R}(A^*)$ and $\mathcal{N}(A)$ is equal to n. ■

This theorem is an example of how a matrix may split a vector space into subspaces. Statement 1 says that the vector space $\mathbb{C}^m$ may be thought of as the direct sum of two subspaces generated by a matrix A. Later on, we shall see how to split a vector space into a direct sum of subspaces with additional features. An arbitrary matrix A is used in the theorem and, hence, the statements contain A and A^*. When A is Hermitian, the statements simplify to a single statement:

$$\mathcal{R}(A) \oplus \mathcal{N}(A) = \mathbb{C}^n$$

Thus, the range space and null space of an Hermitian matrix are orthogonal complements of each other.

The theorem has a geometric flavor and has many consequences. For example, let r be the rank of $A \in \mathbb{C}^{m\times n}$. Statement 1 and the definition of rank give:

$$r + \dim \mathcal{N}(A^*) = m \Rightarrow r \leq m$$

which shows that the rank of a matrix is less than or equal to the number of rows. We had mentioned earlier that, as a consequence of the interpretation of range as the column span, rank is less than or equal to the number of columns.

Proposition 3.5.1 (Some consequences of Theorem 3.5.1) *Let $A \in \mathbb{C}^{m\times n}$ and $B \in \mathbb{C}^{n\times k}$ be of ranks r_A and r_B respectively. The following statements are true.*

1. $r_A \leq \min\{m, n\}$
2. *r_A is equal to the rank of A^**
3. *r_A is equal to the total number of linearly independent rows (or columns)*
4. *r_A is equal to the rank of AA^**
5. *Rank of AB is less than or equal to the minimum of r_A and r_B* ■

We already mentioned the simplification that occurs when A is Hermitian. Even in the general case of square matrices, range space and null space have a curious property. Let A be a square matrix and $x \in \mathcal{R}(A)$. Then, Ax is also in the $\mathcal{R}(A)$ because after all the definition of range of a A is the set of all vectors of the form Ax. This property called *invariance under* A leads to the theory of invariant subspaces.

Definition 3.5.1 (Invariant subspace) *Let $A \in \mathbb{C}^{n\times n}$. A subspace $\mathcal{S}$ of $\mathbb{C}^n$ with the property $A\mathcal{S} \subset \mathcal{S}$, that is,*

$$Ax \in \mathcal{S} \quad \textit{whenever } x \in \mathcal{S}$$

is called an A-invariant subspace. ⋆

Example 3.5.1 (Eigen-subspace) *Let $A \in \mathbb{C}^{n\times n}$ and λ be an eigenvalue of A. The eigen-subspace of A associated with the eigenvalue λ:*

$$\begin{aligned} E_\lambda &= \mathbf{Span}\left(\textit{linearly independent eigenvectors associated with } \lambda\right) \\ &= \{x : (\lambda I - A)\, x = 0\} \end{aligned}$$

is an A-invariant subspace.

Suppose that $\{\lambda_i\}_{i=1}^k$ are eigenvalues of A with the eigensubspaces E_{λ_i}. Then, the direct sum:

$$\bigoplus_{i=1}^{k} E_{\lambda_i}$$

is also an A-invariant subspace. △